

Information-theoretically Secure Erasure Codes for Distributed Storage

Nihar B. Shah, K. V. Rashmi, Kannan Ramchandran, *Fellow, IEEE*, and P. Vijay Kumar, *Fellow, IEEE*

Abstract

Repair operations in distributed storage systems potentially expose the data to malicious acts of passive eavesdroppers or active adversaries, which can be detrimental to the security of the system. This paper presents erasure codes and repair algorithms that ensure security of the data in the presence of passive eavesdroppers and active adversaries, while maintaining high availability, reliability and efficiency in the system. Our codes are optimal in that they meet previously proposed lower bounds on the storage, network-bandwidth, and reliability requirements for a wide range of system parameters. Our results thus establish the capacity of such systems. Our codes for security from active adversaries provide an additional appealing feature of ‘on-demand security’ where the desired level of security can be chosen separately for each instance of repair, and our algorithms remain optimal simultaneously for all possible levels. The paper also provides necessary and sufficient conditions governing the transformation of any (non-secure) code into one providing on-demand security.

I. INTRODUCTION

The amount of data stored in large-scale, distributed storage-systems such as data centers is increasing exponentially. In order to scale massively at low costs, data-centers employ inexpensive commodity hardware. As these components are prone to failure [3], [4], the system must possess enough redundancy to ensure that data remains reliable and available in the face of these failures. One means of introducing redundancy is via replication. However, replication is inefficient with respect to storage space utilization, and thus in order to scale economically, data centers are increasingly turning to the use of erasure coding as a far more efficient option [5], [6].

Consider a distributed storage system with n storage nodes across which some data (termed the *message*) is to be stored. Each of these n nodes stores only a fraction of the data. In order to provide reliability and availability, the erasure codes considered in this paper ensure that a user (termed a *data collector*) must be able to recover the message from the data stored in *any* k of the n nodes. This property is called *data reconstruction* or simply *reconstruction*. The reconstruction property provides the storage system a capability of tolerating failures of *any* $(n - k)$ of the n nodes.

Upon failure of a storage node, a *replacement* node is designated to store the data that was stored previously in the failed node. The replacement node recovers this data by downloading (a part of) the data stored in the remaining nodes. This is termed a *node-repair* or simply a *repair* operation. Traditional erasure codes typically

Nihar B. Shah, K. V. Rashmi and Kannan Ramchandran are with the Dept. of EECS, University of California, Berkeley, CA 94703, USA. Email: {nihar,rashmi,kannan}@eecs.berkeley.edu. P. Vijay Kumar is with the Dept. of ECE, Indian Institute Of Science, Bangalore, India. Email: vijay@ece.iisc.ernet.in. P. Vijay Kumar is also an adjunct faculty member of the Electrical Engineering Systems Department at the University of Southern California, Los Angeles, CA 90089-2565.

This paper was presented in part at Globecom 2011 [1] and in part at ISIT 2012 [2].

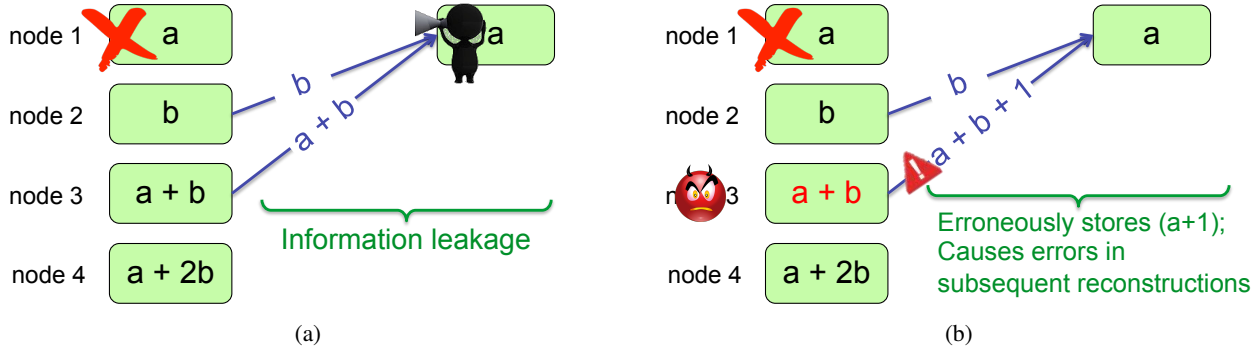


Fig. 1: Compromise of security when using a Reed-Solomon code due to (a) passive eavesdroppers, and (b) active adversaries.

handle repair by first downloading and reconstructing the entire message and then extracting the required data from it. Such an operation is quite wasteful of the network resources [4], [7], [8], and several recent works [7]–[28] propose new erasure codes and repair algorithms addressing this issue.

Security is an important aspect of distributed storage systems, and this is underscored by the many incidents of data compromise in the recent past (e.g., [29]–[31]). This paper considers the problem of designing codes and algorithms for distributed storage that ensure security in addition to maintaining other properties such as reliability, availability and efficiency. We consider the information-theoretic notion of security (and not the computational notion), making no assumption about the computational power of the adversary.

In designing a secure distributed storage system, special attention must be given to the repair operations since they can pose security hazards. For instance, in a system using traditional erasure codes, the repair of a failed node would require the download of the entire data at the replacement node. An eavesdropper tapping onto the replacement node can thus obtain the entire data. Alternatively, during repair of a failed node, an active adversary that has captured one or more of the remaining nodes may pass corrupt data during the repair operation. Such an action would corrupt the data stored in the replacement node, and these errors may then propagate across the entire system during subsequent repairs of other nodes. This makes repair operations a serious security threat, and motivates the goal of this paper.

Fig. 1 illustrates the problem of security under repair dynamics using a toy example of a Reed-Solomon code with $n = 4$ and $k = 2$. The code stores the message $\{a, b\}$ across four nodes, where both a and b belong to the finite field \mathbb{F}_5 in this example. It is easy to verify that the entire message can be recovered by downloading the data stored in any $k = 2$ of the four nodes. On failure of any node, the replacement node connects to any two other nodes, downloads all data stored in them, from which the replacement node recovers the data stored in the node prior to failure. Fig. 1a illustrates a setting with passive droppers. Specifically, consider an eavesdropper who can read all the data stored in node 1. Under the Reed-Solomon code, the eavesdropper can gain access to the symbol a . Furthermore, if the eavesdropper was also able to listen to the data passed for any repair operation, then it gains access to the entire message. Fig. 1b illustrates a setting with active adversaries. Suppose an active adversary gains access to one of the nodes, say node 3, and suppose the repair of node 1 is performed by connecting to nodes 2 and 3. Then, as seen in the figure, the adversary can pass malicious data in the repair process (for instance, $(a + b + 1)$ instead of $(a + b)$), making the replacement node $(a + 1)$ believing it is actually storing a . This error, in turn, propagates during further repairs, and also sabotages all subsequent reconstruction operations involving node 1.

In this paper, we model the distributed storage system based on the *regenerating codes* model introduced by Dimakis et al. [32]. In addition to the parameters n and k introduced earlier, this model has a third parameter d . Upon failure of a node, the replacement node may connect to *any* d of the remaining nodes, and should be able to recover the data that was stored previously in the failed node by downloading a minimal amount of data from these d nodes. It was shown in [32] that there is a tradeoff between the total amount of data stored per node and that downloaded for repair of a failed node, and this is described in greater detail Section II-C. The two extreme regimes in the tradeoff are the minimum storage regenerating (MSR) and the minimum bandwidth regenerating (MBR) regimes. The MSR regime aims to minimize the amount of storage required per node, and for this amount of storage, minimizes the amount of download. On the other hand, the MBR regime aims to keep the amount of download to an absolute minimum, and for this amount of download, minimizes the amount of storage.

The initial work [32] on regenerating codes considers the repair to be ‘functional’, wherein a replacement node need not be identical to the failed node, but should only satisfy further reconstructions and repairs. A strictly stronger and practical requirement is that of ‘exact’ repair, where the replacement node must obtain and store data *identical* to that in the failed node. Throughout this paper, we consider only exact repair.

The threat model we consider for security in this paper is an extension of the threat model proposed in [33] for the regenerating codes setup. Two classes of threats are considered:

- *Security from passive eavesdroppers*: This threat class involves preventing leakage of any information about the data to passive eavesdroppers who may gain access to a subset of the storage nodes. By ‘passive’ we mean that the eavesdropper may read and store any data it gains access to, but does not corrupt any data. In our threat model, two parameters ℓ and m determine the level of security to be provided: the goal is to ensure that a passive eavesdropper having access to the data *stored* in any ℓ of the nodes and additionally to the *data passed for repair* in any m of these ℓ nodes, gains zero information about the message.
- *Security from active (malicious) adversaries*: This threat class involves ensuring that the presence of an adversary who gains access to a subset of the storage nodes and wishes to actively corrupt the data does not affect the operations in the system. The active adversary may corrupt the data stored in a subset of storage nodes and also pass erroneous data to other nodes during their repair operations. In particular, for some parameter p determining the level of security to be provided, the goal is to guard against p nodes being corrupted by an active adversary. In such a scenario, if a replacement node (or a data-collector) connects to any of the p corrupt nodes, then these corrupt nodes could pass arbitrary data in the repair (or data-reconstruction) operation. The goal here is to successfully accomplish the node-repair and data-reconstruction operations even in the presence of such an attack.

In this paper, we present explicit codes for distributed storage at MSR and MBR points that ensure security of the data from passive eavesdroppers and/or active adversaries. The secure codes at the MBR point are applicable for all values of the parameters $[n, k, d]$, and the secure codes at the MSR point are applicable for all values of the parameters satisfying $d \geq 2k - 2$. The secure codes presented have two attractive features:

- *Optimality*: Our codes for security from active adversaries and from passive eavesdroppers at the MBR point are optimal for all values of the parameters. The codes for security from active adversaries at the MSR point is optimal for all parameters satisfying $d \geq (2k - 2)$. The codes for security from passive eavesdroppers at the MSR point are optimal for all parameters satisfying $d \geq (2k - 2)$ with $m \leq 1$. The codes presented in this paper, thus, establish the secure *capacity* of such systems.
- *On-demand security from active adversaries*: Under our codes, when dealing with active adversaries,

the protection level p need not be fixed a priori, but can be chosen flexibly at “run-time” when any repair or reconstruction operation is being executed. The protection level p can be chosen separately and independently for each instance of repair or reconstruction. This “on-demand” security endows the system with the advantage of not having to preset the security level and associated system resources for the worst case. This is in contrast to the traditional models of information-theoretic security (e.g., [33]) that take a “static” approach towards system design where the price is paid, in terms of the reduction in the size of the message stored, corresponding to the magnitude of the “worst-case” security level required. It turns out, perhaps surprisingly, that *our codes do not require any additional storage* to support this on-demand property, and are *optimal for all values of p* .

The problem at hand is closely related to the problem of secure network coding [34]–[42]. The literature on secure network coding primarily considers a multicast setting where one or more receivers wish to obtain the entire data that was transmitted. The threat models typically consider the scenario of link compromise in the networks. The problem we consider, on the other hand, falls into the harder *non-multicast* setting [34], and furthermore, needs to handle the harder case of adversaries or eavesdroppers being able to *capture nodes* in the network [40]. The results of this paper thus establish the capacity of a class of non-multicast networks in the presence of active adversaries or passive eavesdroppers having the ability to compromise nodes. Interestingly, the capacity-achieving codes that we propose in this paper are linear, deterministic and explicit.

While the primary focus of this paper is information-theoretic security, the results of this paper also are also relevant to other applications such as correcting network errors and erasures and for reducing latency of “degraded reads” in data centers. These applications are discussed in more detail in Section II-D.

The rest of the paper is organized as follows. Section II formalizes the system model, and describes our approach towards the code construction problem. Some additional applications of our codes are also discussed in this section. Section III describes related literature. The codes presented in this paper build upon our previously proposed *product-matrix framework* [11], and this framework is overviewed in Section IV. Sections V and Section VI present explicit constructions of MBR codes that provide security from active adversaries and from passive eavesdroppers respectively. Section VII and Section VIII present explicit construction of MSR codes that provide security from active adversaries and from passive eavesdroppers respectively. Section IX provides necessary and sufficient conditions governing the transformation of any (non-secure) erasure code into one providing on-demand security. Section X presents a concluding discussion.

II. SYSTEM MODEL, OUR APPROACH, AND SUMMARY OF THE RESULTS

We will first describe the system model in the absence of security requirements (the ‘regenerating codes’ model), following which we describe the extension that incorporates the provision of security. Alongside, we also describe our approach towards constructing secure distributed storage codes, and a summary of the results of this paper.

A. Regenerating Codes Model (in the absence of security requirements)

Under the regenerating codes model [32], the system comprises of n storage nodes, across which data comprising of B symbols is to be stored. This set of B symbols is called the *message*, and each of these symbols is assumed to belong to a finite field \mathbb{F}_q of size q . Each of the n storage nodes has a capacity of storing α symbols. The data is to be stored such that a user (termed *data-collector*) can recover the entire message by downloading the data stored in *any* k of these n nodes. This process is termed *data-reconstruction*,

or simply *reconstruction*. It follows that any system satisfying the reconstruction property can tolerate the failure of *any* $(n - k)$ storage nodes without losing any data.

Now let us consider the repair operation. When a storage node fails, it is replaced by another node, called the replacement node, that stores exactly the same data as the failed node. The regenerating codes model contains two additional parameters, d and β , that are associated to the repair of failed nodes. The replacement node is permitted to connect to *any* d ($\geq k$) nodes out of the remaining $(n - 1)$ nodes while downloading β ($\leq \alpha$) symbols from each node. This set of d nodes are termed the *helper nodes* for this instance of repair. From the set of $d\beta$ symbols thus obtained, the replacement node must recover the α symbols that were stored in the failed node. The total amount $d\beta$ of data downloaded for repair purposes is termed the *repair bandwidth*.

It is shown in [32] that the parameters associated with a regenerating code must necessarily satisfy the bound

$$B \leq \sum_{i=0}^{k-1} \min(\alpha, (d - i)\beta) . \quad (1)$$

Since both storage and bandwidth come at a cost, it is naturally desirable to minimize both α as well as $d\beta$, and try to achieve the bound (1) with equality, i.e.,

$$B = \sum_{i=0}^{k-1} \min(\alpha, (d - i)\beta) . \quad (2)$$

It can be deduced (see [32]) that achieving (2), for fixed values of B and $[n, k, d]$, leads to a tradeoff between the storage space α and the repair-bandwidth $d\beta$.¹

Two important regimes of this tradeoff are its extremities, termed the minimum storage regenerating (MSR) and minimum bandwidth regenerating (MBR) regimes, described below.

Minimum Bandwidth Regenerating (MBR) regime: The MBR regime entails the minimum possible repair-bandwidth: the amount of data downloaded by a replacement node is no more than the amount of data that was stored in the failed node. To arrive at the MBR regime, one must first minimize the bandwidth $d\beta$ in the tradeoff (2), and then obtain the minimum storage α for this value of repair-bandwidth. Clearly, a replacement node must download at least as much data as what was stored in the failed node. Moreover, it is shown in [43] that for optimal performance, every storage node must utilize its full storage capacity of α symbols. Thus, the minimum bandwidth required is $d\beta = \alpha$. To satisfy (2) with this value of $d\beta$, the parameters associated to the MBR regime must satisfy

$$\alpha = \frac{B}{k} + \frac{(k - 1)\beta}{2} , \quad d\beta = \alpha . \quad (3)$$

Minimum Storage Regenerating (MSR) regime: The MSR regime allows for the minimum possible storage capacity: the requirement of being able to recover all B message symbols from any k nodes mandates the storage per node to be at least $\frac{B}{k}$, and the MSR regime requires attainment of this minimum, i.e., has $\alpha = \frac{B}{k}$. Substituting this value in (2), it follows that the parameters associated to the MSR regime are

$$\alpha = \frac{B}{k} , \quad d\beta = \alpha + (k - 1)\beta . \quad (4)$$

Any code operating in the MSR regime is thus Maximum-Distance-Separable (MDS).

The two extremities of the tradeoff curve have been well-studied in the literature, and there exist several

¹In addition to achieving (2), the parameters must also satisfy the following minimality condition: a reduction in either α or β must result in a violation of (1).

explicit codes [9]–[18] operating in these regimes and achieving (2) for exact repair. On the other hand, it is shown in [12] that at essentially all other points on the tradeoff curve, there *cannot* exist any codes that satisfy (2) for exact repair. Tighter outer bounds have been proposed for these intermediate points on the tradeoff curve [44]–[47].

Having described the basic setup of regenerating codes (in the absence of security), we will now present a description of the setting where security from passive eavesdroppers and/or active adversaries is required.

B. Security from Passive Eavesdroppers

1) Threat Model and Upper Bounds: We consider a threat model that extends the model introduced in [33]. under the model we consider, a passive eavesdropper may gain access to the data stored in a subset of the storage nodes, and possibly also to the data downloaded during repair of some of these nodes. We assume that the eavesdropper may possess unbounded computational power and has complete knowledge of the system protocol. It is required that such an eavesdropper should not be able to obtain any knowledge (in the information-theoretic sense) about the message.

An eavesdropper can gain read-access to the data stored in any arbitrary set of at most ℓ ($< k$) storage nodes. In addition, the eavesdropper may also listen to the data downloaded during (any number of) repair operations of some arbitrary subset of m ($\leq \ell$) of these ℓ nodes. We call such an eavesdropper an $\{\ell, m\}$ -eavesdropper, and term the security of data from such a passive eavesdropper as $\{\ell, m\}$ -security from the passive eavesdropper. Note that since the data passed to a node for repair also contains the data it will finally store, the setting described above can be equivalently stated in the following manner:

Definition 1 ($\{\ell, m\}$ -security from the passive eavesdropper): A code provides $\{\ell, m\}$ -security from the passive eavesdropper if a passive eavesdropper gaining access to only the data stored in any $(\ell - m)$ nodes, and the data passed for repair to any m nodes, obtains zero information about the message.

As an example of this model, consider a peer-to-peer storage system. The m nodes described above may represent nodes that are in a network belonging to the eavesdropper, thereby allowing the eavesdropper to listen to all the data downloaded when these m nodes undergo (possibly multiple) failures and repairs across time. On the other hand, the $(\ell - m)$ nodes may represent the nodes which may be exposed only momentarily, allowing the eavesdropper access to only the data stored.

We now describe an upper bound on the size of the message that can be stored for any code under this system model. In [33], Pawar et al. consider the setting of $m = \ell$ and provide an upper bound on the number of message symbols B^* that can be stored in the information-theoretically secure system. It can be shown easily that even under our extended setting of having $m \leq \ell$, the proof of this upper bound (Theorem 1 of [33]) works for any value of m ($\leq \ell$). Using the bound from [33] for this setting, it follows that the parameters associated to any regenerating code providing $\{\ell, m\}$ -security from passive eavesdroppers must necessarily satisfy

$$B^* \leq \sum_{i=\ell}^{k-1} \min(\alpha, (d-i)\beta) . \quad (5)$$

Specializing to the MBR regime of $\alpha = d\beta$, this bound becomes

$$B^* \leq \left(kd - \frac{k(k-1)}{2} \right) \beta - \left(\ell d - \frac{\ell(\ell-1)}{2} \right) \beta . \quad (6)$$

For the MSR regime, the bound in (5) turns out to be loose. Subsequent to the conference publication [1]

of our code constructions for security from passive eavesdroppers, in [48], the authors have provided tighter upper bounds for the MSR regime:

$$B^* \leq (k - \ell)(\alpha - \beta) . \quad (7)$$

2) *Summary of Our Results for Passive Eavesdroppers:* This paper provides explicit constructions for:

- MBR codes for all parameters $[n, k, d]$, $\{\ell, m\}$ providing information-theoretic security from passive eavesdroppers. These codes are optimal for all values of the parameters.
- MSR codes for all parameters $[n, k, d \geq 2k - 2]$, $\{\ell, m\}$ providing information-theoretic security from passive eavesdroppers. These codes are optimal whenever $m \leq 1$.

The MBR code is optimal by virtue of meeting the outer bound (6) (derived in [33]) with equality. The MSR code is optimal for $m \leq 1$ by virtue of meeting (7) (derived in [48]).

Our codes thus establish the secure capacity of distributed storage systems for the case of passive eavesdroppers for all parameters at the MBR regime and for all $[n, k, d \geq 2k - 2]$ and $m \leq 1$ at the MSR regime.

We take the following approach for constructing the codes. To construct a secure code for a given $[n, k, d]$, we choose a product-matrix (PM) code with identical values of system parameters $[n, k, d]$. In the input to the PM code (without secrecy constraints), we replace a carefully chosen set of

$$R = B - B^* \quad (8)$$

message symbols with R random symbols, where B is as obtained by setting equality in (3) for the MBR case and in (4) for the MSR case, and B^* is the number of message that the code can store under the specified security requirement. Each of these random symbols are chosen uniformly and independently from \mathbb{F}_q , and are also independent of the message symbols. We then prove that this construction ensures that no $\{\ell, m\}$ -eavesdropper can obtain any information about the message, thus ensuring $\{\ell, m\}$ -security from passive eavesdroppers.

3) *Example:* We illustrate our code construction with a toy example of an MBR code with $[n = 3, k = 2, d = 2]$ providing $\{\ell = 1, m = 1\}$ -security from passive eavesdroppers. The code is shown in Fig. 2. The message $\{a\}$ is encoded and stored across $n = 3$ nodes in a manner that it can be recovered from any $k = 2$ of the nodes. The alphabet of operation is the finite field \mathbb{F}_3 . Symbols r_1 and r_2 are drawn uniformly at random from \mathbb{F}_3 . A failed node is repaired by downloading one symbol each from any $d = 2$ nodes. A passive eavesdropper gaining access to the data stored in any $\ell = 1$ node and also to the data passed to that node during any of its repair operations ($m = 1$) gets zero information about the message. For instance, the repair of node 2 is shown in the figure, where one can see that the eavesdropper gains no information about a from the data stored in or passed to any one node. The code is *optimal* meaning that the amount of download required for any repair is the minimum possible, and furthermore, the amount of storage required in this setting is also minimum.

C. Security from Active Adversaries

1) *Threat Model and Upper Bounds:* We consider the threat model wherein one or more nodes may be compromised to an active adversary. The adversary is assumed to be computationally unbounded and knows the protocol followed by the system. A node that is compromised to an active adversary may send arbitrarily corrupted data during any data-reconstruction or node-repair operation. It is required that the data-reconstruction operations complete successfully (without any errors) despite the presence of these active adversaries. This threat

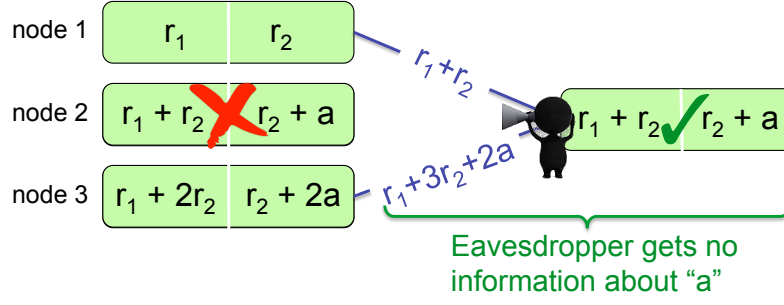


Fig. 2: MBR code with $[n = 3, k = 2, d = 2]$ providing optimal $\{\ell = 1, m = 1\}$ -security from passive eavesdroppers. Repair of node 2 is shown, and one can see that the eavesdropper gains no information about a from the data stored in or passed to any single node.

model was introduced in [33]. Note that this threat model does not require node repairs to be performed without errors, and only requires data-reconstruction operations to be performed without errors.

The upper bound on the size of the message that can be securely stored under the above described threat model with at most p compromised nodes, as derived in [33] is

$$B \leq \sum_{i=2p}^{k-1} \min(\alpha, (d-i)\beta) . \quad (9)$$

Our codes achieve this bound and also ensure that every repair operation is also without corruption.

Along the lines of traditional models of information-theoretic security, [33] takes a “static” approach to secure-code construction, wherein the size of the message stored is reduced by an amount proportional to the magnitude “ p ” of the required security level. In particular, the value of p chosen is typically based on worst-case estimation, which leads to considerable wastage of system resources during the normal operations. Furthermore, in the event that the number of nodes compromised is greater than that anticipated, the security level cannot be increased one-the-fly and hence security can no longer be guaranteed.

Our approach to secure data from active adversaries is different from the conventional approach that we just described. Under our approach, the message is encoded and stored independent of any security requirements, and all the security requirements are handled “on-demand” at the decoding stage by downloading additional data. Let us assume that at some point in time we estimate a requirement of protection from the compromise of any p out of the n nodes in the network. Under our approach, this value of p is *not* fixed at the time of encoding of the data, but possible corruptions are corrected by downloading additional data during the reconstruction or repair processes. Now, during a particular instance of reconstruction or repair, the amount of additional data downloaded is a function of the desired protection level p , and this protection level can be chosen independently for each reconstruction and repair operation. This provides the advantage of not having to allocate resources for the worst case. This approach also enables the system to function with dynamic protection levels, and hence there is no need to estimate and fix the parameter p beforehand: the system administrator is free to choose the desired level of security at runtime.

In our approach, the additional data downloaded for correcting possible corruptions is obtained by allowing a (greater) connectivity of $\Delta (\geq d)$ nodes during repair and $\kappa (\geq k)$ nodes during reconstruction. The parameters

Δ and κ depend on the desired protection level p as

$$\Delta = d + 2p \quad (10a)$$

$$\kappa = k + 2p . \quad (10b)$$

Under this notation the outer bound of [33] translates to

$$B \leq \sum_{i=2p}^{\kappa-1} \min(\alpha, (\Delta - i)\beta) . \quad (11)$$

Observe that substituting the values of κ and Δ in terms of k and d from (10) in (11) results in (1). It follows that achieving (11) with equality for all p is equivalent to achieving (10) and (1) with equality for all p . In particular, it amounts to achieving (10) along with (3) for the MBR regime and (4) for the MSR regime.

Our goal is to construct codes which can secure the data from active adversaries for any value of the desired level of security p that can be chosen separately at each instance of repair or reconstruction. We also require that, in addition, the codes be optimal simultaneously for all values of p . We term this *on-demand security*.

Definition 2 (On-demand security): A code is said to provide on-demand security if for *any* arbitrary (feasible) choice of p made at the time of a repair (or reconstruction) operation, the repair (or reconstruction) algorithm can detect and correct all corruptions in the presence of up to p compromised nodes.

Note that by all feasible values of p , we mean all values of p for which the additional connectivity requirement described in (10) can be satisfied (i.e., for all p such that $(d + 2p) \leq (n - 1)$).

Remark 1 (On-demand detection of errors): Note that the definition of on-demand security also requires *detection* of corruption, when up to p of the nodes are compromised during any repair (or reconstruction) operation. The corruption detection can be performed by connecting to $(d+p)$ (or $(k+p)$) nodes and downloading β (or α) symbols from each of them during any repair (or reconstruction) operation as opposed to $(d + 2p)$ (or $(k + 2p)$) nodes for correction of corruption. The error detection properties of the codes presented in this paper follow from their error correction properties in a straightforward manner using the basic relationship between detection and correction of errors in block codes [49]. We hence focus only on the error correction in the rest of the paper, and the error detection properties follow as implicit corollaries.

In our codes providing on-demand security, the parameters $[n, k, d]$ and $\{\beta, \alpha, B\}$ remain fixed, while the values of Δ and κ may vary during different instances of reconstruction or repair operations depending on the security level p desired at that time. Furthermore, our codes achieve the bound (11) with equality for all values of p , and are hence optimal.

2) *Summary of Our Results for Active Adversaries:* This paper presents explicit constructions for:

- MBR codes for all parameters $[n, k, d]$ providing on-demand security from active adversaries. The codes are optimal for all values of the parameters.
- MSR codes for all parameters $[n, k, d \geq 2k - 2]$ providing on-demand security from active adversaries. The codes are optimal for all values of the parameters.

The optimality of the codes follows from the fact that the MBR codes meet (3) and (10) for all values of the parameters and the MSR codes meet (4) and (10) for all values of the parameters covered. Our results thus establish the secure capacity of distributed storage systems for the case of active adversaries for the aforementioned parameter regimes.

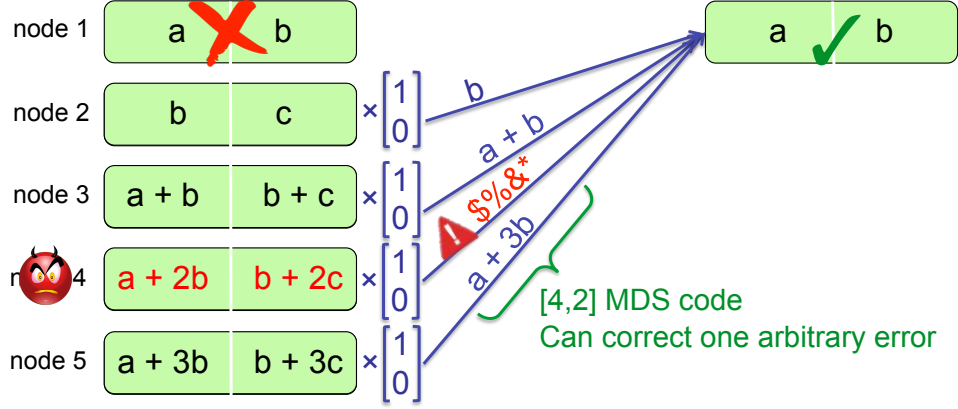


Fig. 3: MBR code with $[n = 5, k = 2, d = 2]$ providing optimal and on-demand security. Repair of node 1 is shown when security from at most $p = 1$ compromised node is required. We can see that the repair operation proceeds successfully in spite of the presence of a compromised node that sends arbitrarily corrupted data.

3) *Example:* We illustrate our code construction with an example of an MBR code with $[n = 5, k = 2, d = 2]$ providing optimal and on-demand security. The code is shown in Fig. 3. The message $\{a, b, c\}$ is encoded and stored across $n = 5$ nodes in a manner that it can be recovered from any $k = 2$ of the nodes. The alphabet of operation is the finite field \mathbb{F}_5 . In the scenario that a repair operation needs to be made secure from the compromise of any $p = 1$ node, the replacement node connects to any 4 nodes and downloads one symbol from each as shown in the figure. Here, even in the presence of one arbitrary corruption (node 4 in the example), the desired symbols a and b are decoded correctly. When no security is required ($p = 0$), a failed node is repaired by connecting to any $d = 2$ nodes and downloading one symbol from each of them: the symbol passed by any node is identical to what it would have passed in the $p = 1$ case. This code is optimal and provides on-demand security: a repair operation may choose the desired level of protection p at runtime and the amount of network-bandwidth consumed in either case is the minimum possible, and furthermore, the amount of storage required to support this amount of download is also minimum.

D. Other Applications

While the codes constructed in this paper are presented from the perspective of providing security in a distributed storage system, they can also be employed for other relevant applications.

1) *Handling Packet Errors and Erasures in the Network:* During a node-repair or a data-reconstruction operation, due to noise in the network, one or more packets downloaded may contain (arbitrary) errors. This is equivalent to a situation in the security scenario in which a set of nodes that are compromised to an active adversary may transmit corrupt packets. Our codes for on-demand security from active adversaries can alternatively be employed for handling packet-errors and erasures in the network. The reader is referred to a conference version [2] of this paper for more details on this application.

The codes for security can correct for packet erasures as well. In this case, a node-repair (or data-reconstruction) operation guarding against p errors and p' erasures is carried out by downloading β (or α) symbols each from any arbitrary $(d + 2p + p')$ (or $k + 2p + p'$) nodes. An erasure of some p' of these packets leaves us with data from some arbitrary subset of $(d + 2p)$ (or $(k + 2p)$) nodes, p of which could be in error. This resultant scenario is identical to the setting of Section II-C, which is addressed by our codes. The properties of on-demand security

and optimality in the amount of download are also retained in the case of erasures.

2) *Reducing Latency of Degraded Reads in Data Centers*: The property of on-demand security can also aid in the reducing the latency of *degraded reads* in data centers. A degraded read (also called *online repair*) is an operation which is executed when a request comes in for data stored in a node that is busy or has failed. Under this operation, the request is met by downloading and recovering the requested data from the remaining nodes via a repair operation. Degraded reads are typically latency-sensitive, i.e., they must be served quickly.

If the underlying code provides on-demand security, then it allows for the degraded read operation to be executed faster in the following manner. Let f denote the busy/failed node that stores the requested data. Now, one could use the repair property of the erasure code to recover the required data from any d other nodes. However, due to various sources of randomness in the system such as congestion at the nodes and delay in transmission, the replies from these d nodes will typically arrive at different times. As a result, the net latency will be the maximum of the response-times of these d nodes. On the other hand, a code providing on-demand security allows a reduction in latency by means of sending ‘redundant-requests’ [50]–[54] as follows. For some choice of a parameter r ($d \leq r \leq n$), connect to any r nodes and ask each of these nodes to aid in the ‘repair’ of node f . Since the code provides on-demand security, one can recover the data of node f (and hence serve the request) from the data obtained from the *first d nodes that reply* (the requests to the remaining nodes are then canceled). This lowers the latency of serving this request to the response-time of the d th fastest among the r ($d \leq r \leq n$) helper nodes.

E. Notational Conventions

A vector will be treated as a column vector by default, and a row vector will be written as the transpose of the corresponding column vector. The transpose of a vector or a matrix will be denoted by a superscript T . The term ‘randomly drawn’ will mean ‘drawn uniformly at random’.

III. RELATED LITERATURE

A. Codes for Distributed Storage

The ‘regenerating codes’ model introduced in [32] considers optimizing two important resources: the storage capacity required per node, and the repair-bandwidth. It was shown in [32] that there exists a tradeoff between these two resources, and lower bounds on their requirements were derived. Subsequent to the work of [32], several explicit codes [9]–[12], [14], [16], [18], [55]–[58] were constructed for the MSR and the MBR regimes of regenerating codes, many of which meet these bounds. Furthermore, it was shown in [12], [44], [59] that the bounds are loose at essentially all points in the interior of the tradeoff curve.

A class of explicit codes that are optimal in terms of the storage and bandwidth requirements are the product-matrix codes proposed in [11]. The results of this paper are based on the product-matrix codes, and exploit certain unique features of the underlying product-matrix framework. A detailed description of the product-matrix codes is provided in Section IV.

The requirement of security in the presence of repair dynamics was first considered in [33]. In [33], the authors provide lower bounds on the storage and bandwidth requirements under such a setting. Secure codes for the MBR regime with $d = n - 1$ are also provided, which are based on the repair-by-transfer code of [9], [12]. The codes of [33] providing security from active adversaries, also for the MBR regime with $d = n - 1$, allow propagation of the errors during the repair operation, and treat the resulting errors only at the reconstruction stage.

Such a property may not be desirable in practice since the system administrators may not approve of allowing errors to linger around and propagate through the system. The secure codes presented in the present paper are applicable to any choice of the system parameters $[n, k, d]$ in the MBR regime, and any $[n, k, d \geq 2k - 2]$ in the MSR regime. Furthermore, our codes ensure the correction of errors during every individual repair and the reconstruction operation. While this additional requirement clearly makes the system more practical, we show that surprisingly, this more stringent requirement does *not* necessitate any additional storage or bandwidth requirements.

Subsequent to the initial presentation of our codes in [1], there have been several other works on information theoretic security in distributed storage systems with repair considerations ([48], [60]–[63]). In [48], the authors consider security from passive eavesdroppers and present outer bounds for the MSR setting that are tighter than those proposed in [33]. The authors also present a secure MSR code construction for $d = n - 1$ which employs Zigzag codes [15] along with a maximum rank distance code. These codes are optimal for $d = n - 1$, $m \leq 2$. The secrecy capacity at the MSR regime for $m > 2$ remains open. In [61], the authors provide a tighter upper bound on the file size that can be stored in a system secure from passive eavesdroppers for codes employing linear encoding and decoding at the MSR point. The code construction provided in [48] meets the upper bound for linear codes provided in [61] for $d = n - 1$.

Security from active adversaries is considered in [60]. A different adversarial model, where an active adversary can replace the content of an affected node only once, is considered and bounds and achievable schemes for this setting are provided. The paper also considers the MSR setting providing security from an active adversary who can replace the contents of affected nodes an unbounded number of times, and provides schemes that are optimal for a specific choice of the parameters. In [62], the authors provide information theoretic upper bounds for the cases $n \leq 4$ and $k = d = n - 1$. For these parameters, it is shown that MBR is the one and only regime for which the resource requirements in the system do not increase when, in addition to security from an eavesdropper being able to read data stored in nodes, security from an eavesdropper who can also tap on to the data downloaded during repair is desired. Security in distributed storage systems with heterogeneous storage nodes is considered in [63].

In [64], the authors deal with byzantine fault tolerance by employing product-matrix codes [11]. They use a cyclic redundancy check (CRC) to check the integrity of data during repair and reconstruction, and a feedback scheme to iteratively correct them. However, CRC based schemes are not applicable in the present setting of information-theoretic security since the CRC may also be corrupted by the adversary. The present paper takes a more fundamental look at the problem of handling corruptions in regenerating codes from an information-theoretic perspective.

In [65], the authors derive bounds to determine the secure capacity in a ‘cooperative-repair’ setting [66]. The bounds of [65] show that such an attempt to cooperatively repair may adversely affect the system in the presence of malicious adversaries. In [67] authors study the security of distributed storage systems in the presence of a trusted verifier. In [68], authors propose using polytope codes to address the issue of security against active adversaries.

B. Shamir’s secret sharing

A possible method to ensure information-theoretic security from passive eavesdroppers is to employ Shamir’s secret sharing scheme [69]. Under Shamir’s secret sharing scheme, the data is encoded and stored in a set of n nodes such that the entire data can be recovered from any k nodes, while access to data in any $(k - 1)$

or fewer nodes provides zero information about the data. Now, during repair of a failed node, this scheme requires a download of the entire data to a central location, following which the data in the replacement node is re-encoded. Thus, as in the case of classical erasure codes, the repair operations are inefficient, mandating significant network resources. Furthermore, this central location represents a single point of failure, and security in such a system can be compromised by an eavesdropper who gains access to this location. These critical issues thus necessitate investigation of alternative solutions that account for the routinely performed failure-handling tasks.

C. Secure Network Coding

The literature on secure network coding (e.g., [35], [39], [41], [42]) primarily considers a *multicast* setting in which there is a single source of data and every destination is interested in obtaining *all* the data transmitted by the source. Furthermore, with respect to security from passive eavesdroppers in the multicast setting, only the scenarios where the eavesdropper can access subsets of *links* is well understood in the literature. The problem of secure distributed storage considered in this paper requires handling the case when nodes are compromised. The problem of node-compromise is typically treated as a case of link-compromise by assuming that the eavesdropper gains access to all links that are incident upon the compromised nodes. In our problem, since a node may be repaired by connecting to *any* d nodes, schemes following such an approach cannot have a non-zero rate of transmission in general.

Even in the case of active adversaries, the problem of handling compromise of nodes is much harder [40] than handling compromise of links. Furthermore, our problem falls into the harder setting of non-multicast, whereas the literature on secure network coding primarily considers the multicast setting. The works [37], [38] present results showing guarantees of error correction under random linear network coding under the condition that the subspace obtained at the receiver has a sufficient intersection with the transmitted subspace, a condition that is not guaranteed in our setting. In [36], the authors propose schemes to transmit a message equal to the difference between the largest message that can be sent in the absence of secrecy requirements and a bound on the number of compromised links. In our problem, this difference is almost always 0 or smaller since a compromised node may assist in any number of repairs of any other node. The results of the present paper, on the other hand, establish the capacity of a class of non-multicast networks in the presence of active adversaries or passive eavesdroppers having the ability to compromise nodes. Interestingly, these capacity-achieving codes that we propose in this paper are linear, deterministic and explicit.

The problem of providing information-theoretic secrecy in distributed storage systems is also related to the Wiretap Channel II [70] where an eavesdropper, who is listening to any arbitrary subset of symbols (of fixed size) being transmitted over a noiseless point-to-point channel, should obtain no information about the original message. While schemes providing secrecy in a distributed storage system with only the reconstruction requirement would follow directly from [70], the requirement of addressing node-repair makes the problem non-trivial.

IV. THE PRODUCT-MATRIX FRAMEWORK

The secure codes presented in this paper are based on the *product-matrix* codes [11], and this section describes their underlying framework. These codes operate at both extremities of the storage-bandwidth tradeoff for the parameters:

- 1) MBR: all parameters $[n, k, d]$, and

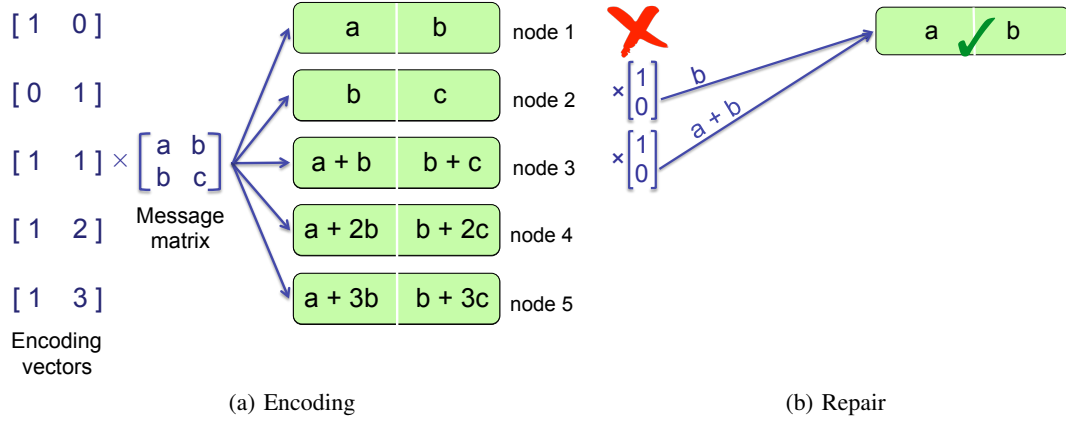


Fig. 4: An example illustrating the product-matrix framework. The encoding of an MBR code for parameters $[n = 5, k = 2, d = 2]$, $(B = 3, \alpha = 2, \beta = 1)$ and $\ell = m = 0$ (i.e., no requirement of security from passive eavesdroppers) is depicted. The field of operation is \mathbb{F}_5 . Also depicted is the repair of the first node with $p = 0$ (i.e., no requirement of security from active adversaries).

2) MSR: all parameters $[n, k, d \geq 2k - 2]$.

Observe from (3) and (4) that for each of these cases, the values of B and α are multiples of β . It follows that given an optimal code for $\beta = 1$, optimal codes for any higher value of β (say, β_0) can be obtained by simply concatenating the $\beta = 1$ code β_0 number of times. This process is known as *striping*. Thus in this paper, without loss of generality, we consider only the case of $\beta = 1$.

Product-matrix codes are represented in terms of an $(n \times \alpha)$ code matrix

$$C = \Psi M, \quad (12)$$

where

- the ‘code matrix’ C is an $(n \times \alpha)$ matrix with its i th row comprising the α symbols stored in node i ($1 \leq i \leq n$),
- the ‘message matrix’ M is a $(d \times \alpha)$ matrix, whose elements comprise all the message symbols, along with some random symbols, arranged in a specific (and possibly redundant) manner, and
- the ‘encoding matrix’ Ψ is a fixed $(n \times d)$ matrix. Denoting the i^{th} row of Ψ by ψ_i^T , the α symbols stored in node i are

$$\psi_i^T M.$$

The d length vector ψ_i^T is termed as the *encoding vector* of node i .

Thus each node stores the inner product of its encoding vector with the message matrix. The specific structures of the matrices M and Ψ vary with the choice of the operating regime (MBR or MSR) and the level of security required, and their design is described in the subsequent sections. Fig. 4 depicts an example of a code constructed under the product-matrix framework.

Systematic Codes: A systematic code is one wherein some set of k nodes together store all the B message symbols in an uncoded form. This set of k nodes are called systematic nodes. In several applications of interest, it may be desired to have the code in a systematic form, since decoding of the data is not required whenever the data-collector connects to the k systematic nodes.

The codes for security from active adversaries that presented in this paper can be constructed in both systematic or non-systematic forms. On the other hand, a code operating under the requirement of information-theoretic security from passive eavesdroppers *cannot* be systematic. This is because, if any part of the message is stored in an uncoded form in any node, an eavesdropper tapping onto that node would obtain a non-zero amount of information about the message.

The following sections employ the product-matrix framework to construct secure MBR and MSR codes.

V. MBR CODES SECURE FROM ACTIVE ADVERSARIES

MBR codes achieve minimum possible download during repair: a replacement node downloads only what it stores, resulting in $d\beta = \alpha$. Recall from Section II-A (see (3)) that in the absence of security requirements, an MBR code must satisfy

$$\alpha = \frac{B}{k} + \frac{(k-1)\beta}{2}, \quad d\beta = \alpha. \quad (13)$$

In this section we present explicit constructions of optimal MBR codes for all parameter values $[n, k, d]$, providing security from active adversaries. As discussed in Section II-C, our code constructions, in addition to being optimal, provide on-demand security. These codes are optimal in the sense that they achieve (11) with equality for every value of the desired security level p .

A. Encoding algorithm

The encoding procedure of the codes for security from active adversaries is identical to that of the product-matrix codes without security requirements [11], and is independent of the desired protection levels. This procedure is briefly described below for the sake of completeness.

We apply the *striping* procedure described in Section IV, and construct codes for $\beta = 1$. Identical encoding, repair and reconstruction operations are performed independently on each stripe of $\beta = 1$. Setting $\beta = 1$ in (13) gives

$$\alpha = \frac{B}{k} + \frac{k-1}{2}, \quad d = \alpha, \quad \beta = 1. \quad (14)$$

Recall from Section IV that the product-matrix codes are represented in terms of an $(n \times \alpha)$ code matrix

$$C = \Psi M, \quad (15)$$

where the i^{th} row of C contains the α symbols stored in node i ($1 \leq i \leq n$). Thus the i^{th} node stores $\psi_i^T M$ where ψ_i^T is the i^{th} row Ψ , and is termed the encoding vector of node i .

The encoding matrix Ψ for the MBR code is of the form

$$\Psi = [\Phi \quad \Sigma], \quad (16)$$

where Φ is an $(n \times k)$ matrix, and Σ is an $(n \times (d-k))$ matrix. The matrices Φ and Σ are chosen such that: (i) any k rows of Φ are linearly independent, and (ii) any d rows of Ψ are linearly independent. These requirements can be met, for example, by choosing Ψ to be either a Cauchy or a Vandermonde matrix [71]. Observe that these properties make Φ the generator matrix of an $[n, k]$ MDS code, and Ψ the generator matrix of an $[n, d]$ MDS code. The choice of the finite field \mathbb{F}_q is restricted (only) by the matrix Ψ : choosing Ψ as a Vandermonde matrix permits any $q \geq n$.

The message matrix M contains the B message symbols arranged in a specific form. As mentioned previously, in our approach, the encoding mechanism is independent of the number of compromised nodes in the system. Hence the number of message symbols is equal to that in the regenerating codes setting with no security requirements (14):

$$B = kd - \frac{k(k-1)}{2} = k(d-k) + \frac{k(k+1)}{2}. \quad (17)$$

The $(d \times d)$ message matrix is designed to be of the form

$$M = \begin{bmatrix} S & V \\ V^T & 0 \end{bmatrix}, \quad (18)$$

where S is a $(k \times k)$ symmetric matrix, and V is a $(k \times (d-k))$ matrix. The B message symbols populate the $\frac{k(k+1)}{2}$ distinct entries of S and the $k(d-k)$ entries of V . Observe that this specific structure makes the message matrix M symmetric.

In the MBR toy example provided in Fig. 4, with $[n = 5, k = 2, d = 2]$ and $(\beta = 1, \alpha = 2, B = 3)$, the (5×2) encoding matrix Ψ and the (2×2) message matrix M take the values:

$$\Psi = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{bmatrix}, \quad M = \begin{bmatrix} a & b \\ b & c \end{bmatrix},$$

with a , b , and c as the three message symbols.

Remark 2 (Systematic code): The PM-MBR code can be made systematic by choosing the encoding matrix Ψ as

$$\Psi = \begin{bmatrix} I_k & 0 \\ \tilde{\Phi} & \tilde{\Sigma} \end{bmatrix}, \quad (19)$$

where I_k is the $(k \times k)$ identity matrix, 0 is a $(k \times (d-k))$ zero matrix, $\tilde{\Phi}$ and $\tilde{\Sigma}$ are matrices of sizes $((n-k) \times k)$ and $((n-k) \times (d-k))$ respectively, such that $\begin{bmatrix} \tilde{\Phi} & \tilde{\Sigma} \end{bmatrix}$ is a Cauchy matrix (or any matrix such that all of its submatrices are of full rank). The example in Fig. 4 depicts a systematic product-matrix MBR code.

This completes the description of the encoding procedure. We now present the decoding algorithms that ensure security from active eavesdroppers.

B. Algorithm for Secure Node Repair

The following theorem presents an explicit algorithm for node-repair in the presence of compromised nodes.

Theorem 1 (Secure node-repair): In the code presented, for any value of parameter p , the repair of a failed node can be secured from the compromise of up to p nodes, by letting the replacement node download $\beta = 1$ symbols each from $(d+2p)$ arbitrary nodes. The amount of data downloaded in this processes is minimum, thus establishing the secure capacity of such systems. The choice of parameter p is arbitrary, and can be different for every instance of repair, thus providing on-demand security.

Proof: Suppose during any instance of node-repair, it is desired to provide security from the compromise of up to p nodes, for some choice of p . Let f denote the failed node and ψ_f be its encoding vector. The $\alpha (= d)$

symbols stored in the failed node are

$$\psi_f^T M .$$

The replacement node then connects to some arbitrary $(d + 2p)$ nodes. The replacement node is required to recover these d symbols by downloading $\beta = 1$ symbol each from the $(d + 2p)$ nodes that it connects to, when at most p of these nodes are compromised. Let \mathcal{J} denote the set of these $(d + 2p)$ nodes. Under our repair algorithm, each of these $(d + 2p)$ nodes pass the inner-product of the α values stored in them with the encoding vector of the failed node f . That is, for every $j \in \mathcal{J}$, node j passes the value of

$$\psi_j^T M \psi_f .$$

This value may be corrupt if node j is compromised to the adversary. Let $\Psi_{\mathcal{J}}$ denote the $((d + 2p) \times d)$ submatrix of Ψ , with its $(d + 2p)$ rows comprising the vectors $\{\psi_j^T\}_{j \in \mathcal{J}}$. Under this notation, one can see that the replacement node has access to the $(d + 2p)$ encoded symbols

$$\Psi_{\mathcal{J}} M \psi_f ,$$

of which at most p symbols may be in error. This set of symbols can equivalently be viewed as an encoding of $M \psi_f$ by the generator matrix $\Psi_{\mathcal{J}}$. To complete the decoding process, we now call upon the MDS property of matrix Ψ . By construction, Ψ is the generator matrix of an $[n, d]$ MDS code. Hence its submatrix $\Psi_{\mathcal{J}}$ is the generator matrix of a $[d + 2p, d]$ MDS code. This implies that the code with $\Psi_{\mathcal{J}}$ as its generator matrix has a minimum distance of $(2p + 1)$. Thus, when at most p of the $(d + 2p)$ symbols $\Psi_{\mathcal{J}} M \psi_f$ are corrupt, the replacement node can apply standard decoding algorithms for MDS codes and recover $M \psi_f$ correctly. Finally, since the message matrix M is symmetric, the d elements comprising $M \psi_f$ can be written as

$$(M \psi_f)^T = \psi_f^T M ,$$

which are precisely the d values that were stored in the failed node.

The repair algorithm described above required the replacement node to connect to some arbitrary $\Delta = (d + 2p)$ nodes, and download $\beta = 1$ symbol from each. The amount of storage and repair-bandwidth consumed meet (17) and (5) with equality, thus establishing optimality. ■

Example 1: Fig. 3 depicts an example of this algorithm where repair of the first node is performed, while providing protection from up to $p = 1$ compromised nodes. The encoding process of the code is illustrated in Fig. 4. For the repair of node 1 as shown in the figure, the replacement node connects to $\Delta = d + 2p = 4$ other nodes and each of these nodes pass an inner-product of their data with $\psi_1 = [1 \ 0]$, i.e., pass their first symbol. The data $\{b, a + b, a + 2b, a + 3b\}$ is an MDS encoding of $\{a, b\}$ allowing for the correction of up to one arbitrary error.

C. Algorithm for Secure Data Reconstruction

The following theorem presents an explicit algorithm for data reconstruction in the presence of compromised nodes.

Theorem 2 (Secure data-reconstruction): In the code presented, a data-collector can recover all the B message symbols by downloading data stored in $(k + 2p)$ arbitrary nodes, in the presence of up to p compromised nodes. The choice of parameter p is arbitrary, and can be different for every instance of data-reconstruction, thus providing on-demand security.

Proof: Suppose during any instance of data-reconstruction, it is desired to provide protection from the compromise of up to p nodes, for some choice of p . The data-collector then connects to $(k + 2p)$ arbitrary nodes. Let \mathcal{I} denote the set of $(k + 2p)$ nodes to which the data-collector connects. Then, from every node $i \in \mathcal{I}$, the data collector downloads the α symbols

$$\psi_i^T M ,$$

which could be corrupt if node i is compromised. Let $\Psi_{\mathcal{I}}$ denote the $((k + 2p) \times d)$ submatrix of Ψ , with its $(k + 2p)$ rows comprising $\{\psi_i^T\}_{i \in \mathcal{I}}$. From the structure of the Ψ (see (16)) we have

$$\Psi_{\mathcal{I}} = \begin{bmatrix} \Phi_{\mathcal{I}} & \Sigma_{\mathcal{I}} \end{bmatrix} , \quad (20)$$

where $\Phi_{\mathcal{I}}$ is a $((k + 2p) \times k)$ matrix and $\Sigma_{\mathcal{I}}$ is a $((k + 2p) \times (d - k))$ matrix. The $(k + 2p)$ rows of $\Phi_{\mathcal{I}}$ and $\Sigma_{\mathcal{I}}$ respectively are comprised of $(k + 2p)$ of the rows of Φ and Σ (indexed by \mathcal{I}). From the specific structure (18) of the message matrix M , we see that the data-collector equivalently has access to the encoded symbols

$$\Psi_{\mathcal{I}} M = \begin{bmatrix} \Phi_{\mathcal{I}} S + \Sigma_{\mathcal{I}} V^T & \Phi_{\mathcal{I}} V \end{bmatrix} .$$

Note that the $(k + 2p)$ rows of the matrix $\Psi_{\mathcal{I}} M$ are obtained from $(k + 2p)$ different helper nodes. Thus, a compromise of up to p of the helper nodes to an active adversary leads to at most p of the rows of $\Psi_{\mathcal{I}} M$ being corrupt.

We now call upon the MDS property of Φ to complete the reconstruction process. Recall that by construction, Φ is the generator matrix of an $[n, k]$ MDS code. Hence its submatrix $\Phi_{\mathcal{I}}$ is the generator matrix of an $[k + 2p, k]$ MDS code. It follows that a code with $\Phi_{\mathcal{I}}$ as its generator matrix has a minimum distance of $(2p + 1)$, and thus has the ability to correct up to p arbitrary errors. Consider the set of symbols $\Phi_{\mathcal{I}} V$ that the data collector has obtained. Observe that any column of the matrix $\Phi_{\mathcal{I}} V$ is an encoding of the corresponding column of V with $\Phi_{\mathcal{I}}$ as the generator matrix. Thus, when at most p of the nodes are compromised, the data-collector can correctly decode each column of $\Phi_{\mathcal{I}} V$ separately using standard algorithms for decoding MDS codes. At the end of this process, the data-collector recovers the matrix V correctly. The next step is to decode S from the downloaded data. Since the value of V is (correctly) known, the data-collector can subtract $\Sigma_{\mathcal{I}} V^T$ from $(\Phi_{\mathcal{I}} S + \Sigma_{\mathcal{I}} V^T)$ to recover $\Phi_{\mathcal{I}} S$. Again, the matrix $\Phi_{\mathcal{I}} S$ is an encoding of the columns of S by the generator matrix $\Phi_{\mathcal{I}}$, which allows for the decoding of S in a manner identical to the decoding of V even in the presence of upto p corrupt rows.

In this setting, one can verify that the amounts of storage and repair-bandwidth consumed meet (9) with equality, thus establishing optimality. \blacksquare

VI. MBR CODES FOR SECURITY FROM PASSIVE EAVESDROPPERS

In this section, we present explicit constructions of MBR codes secure from passive eavesdroppers that support all values of $[n, k, d]$ and all $\{\ell, m\}$. As discussed previously, unlike our codes for security from active adversaries, the codes for security from passive eavesdroppers do not provide on-demand security. As a result, the desired level $\{\ell, m\}$ of security must be chosen apriori at the time of encoding. The MBR regime is governed by the relation $\alpha = d\beta$, and substituting this in (5) gives an upper bound on the amount of data B^* that can be stored in a $\{\ell, m\}$ -secure MBR code as

$$B^* \leq \left(kd - \frac{k(k-1)}{2} \right) \beta - \left(ld - \frac{\ell(\ell-1)}{2} \right) \beta . \quad (21)$$




Fig. 5: Structure of the $(d \times d)$ message matrix M^* for the $\{\ell, m\}$ -secure MBR code. The matrix is symmetric and comprises of blocks of data, random symbols, and zeros. The code is given by $C^* = \Psi^* M^*$ where Ψ^* is the encoding matrix.

The codes presented here achieve this bound with equality

$$B^* = \left(kd - \frac{k(k-1)}{2} \right) \beta - \left(ld - \frac{\ell(\ell-1)}{2} \right) \beta, \quad (22)$$

thus establishing this as the secure capacity under the MBR regime for all values of the parameters $[n, k, d]$ and $\{\ell, m\}$.

When any parameter takes different values in the secure and non-secure versions, we will denote the parameter associated to a code providing security with a superscript asterisk (*) as done in (21) and (22).

A. Encoding

We will now construct an $\{\ell, m\}$ -secure MBR code satisfying (22). We apply the striping procedure described in Section IV, and (without loss of generality) construct codes for the case $\beta = 1$. Setting $\beta = 1$ in (22) gives

$$B^* = \left(kd - \frac{k(k-1)}{2} \right) - \left(ld - \frac{\ell(\ell-1)}{2} \right), \quad \alpha = d, \quad \beta = 1.$$

The product-matrix codes for security from eavesdroppers possess a structure identical to those without security requirements. The difference between the two codes is that in the absence of security requirements, the elements of the message matrix M consist of the message symbols, whereas in the situation when security is needed, a specific set of elements of the message matrix are replaced by random symbols. Thus, to construct an $\{\ell, m\}$ -secure code for any $[n, k, d]$, first consider an $[n, k, d]$ product-matrix MSR code in the absence of security requirements, as constructed in Section V-A. Denote this code by \mathcal{C} , and further, let \mathcal{C}^* denote the $\{\ell, m\}$ -secure code (which we will construct below). The code \mathcal{C} is described by an $(n \times \alpha)$ code matrix $C = \Psi M$ as defined in (15), with the $(d \times d)$ matrix M comprising $B = kd - \frac{k(k-1)}{2}$ message values (14). Following the product-matrix framework, the code \mathcal{C}^* is also described by an $(n \times \alpha)$ code matrix C^* of the form $C^* = \Psi^* M^*$. The matrices Ψ^* and M^* for code \mathcal{C}^* are obtained by modifying the matrices Ψ and M of code \mathcal{C} in a manner described below.

The $(n \times d)$ encoding matrix Ψ^* is required to satisfy the following property in addition to those required by Ψ : when restricted to the first ℓ columns, any ℓ rows are linearly independent. The choice of Ψ^* as a Cauchy or a Vandermonde matrix satisfies this property as well.

The $(d \times d)$ message matrix M^* is obtained by replacing the

$$R = B - B^* = \ell d - \frac{\ell(\ell - 1)}{2} \quad (23)$$

message symbols in the first ℓ rows (and hence first ℓ columns) of the symmetric matrix M by R random symbols. Each random symbol is drawn independently and uniformly from \mathbb{F}_q . The structure of the resulting message matrix is illustrated in Fig. 5.

Finally, node i ($1 \leq i \leq n$) stores the i^{th} row of the $(n \times \alpha)$ code matrix $C^* = \Psi^* M^*$.

The following example illustrates the encoding procedure.

Example 2: Consider the code depicted in Fig. 2 providing security from passive eavesdroppers. This code operates in the MBR regime with parameters $n = 3$, $k = 2$ and $d = 2$, and provides security against eavesdroppers who may gain access to $(\ell = 1, m = 1)$ nodes. Let the alphabet of operation be \mathbb{F}_3 and let $\alpha = 2$ and $\beta = 1$. From (22) we get that the maximum size of the message that can be stored securely in this system is $B^* = 1$ symbol (we could have stored $B = 3$ symbols in the absence of security requirements). Let us denote this solitary message symbol as a . Following (23), we draw $R = 2$ symbols r_1 and r_2 uniformly at random from \mathbb{F}_3 . The encoding vector of any node i ($1 \leq i \leq 3$) under this code is $\psi_i^T = [1 \ i]$. The message matrix is

$$M^* = \begin{bmatrix} r_1 & r_2 \\ r_2 & a \end{bmatrix}.$$

Each node i stores $\psi_i^T M^*$ as shown in Fig. 2.

B. Reconstruction, Repair and Security

The following theorems prove the properties of reconstruction, repair and security.

Theorem 3 (Data-reconstruction and node-repair): In code C^* , a data-collector can recover all the B^* message symbols by downloading the data stored in any k nodes, and a failed node can be repaired by downloading $\beta = 1$ symbol each from any d remaining nodes.

Proof: Treating the random symbols also as message symbols, code C^* becomes identical to C . Thus reconstruction and repair in C^* are identical to that in C . ■

Note that repair and reconstruction can be carried out ensuring security from corruption by active adversaries by employing the explicit algorithms of Section V-B and Section V-C.

Theorem 4 (Information-theoretic security): The code C^* ensures $\{\ell, m\}$ -security, i.e., an eavesdropper having access to the data stored in up to ℓ nodes, and to all the data passed to up to m of these nodes during one or more of their repair operations, gets no information about the message.

Proof: The complete proof is provided in the appendix, and we provide an outline here. Let \mathcal{U} , \mathcal{R} and \mathcal{E} respectively denote random variables corresponding to the message, the random symbols inserted at the encoding stage, and the data available to the eavesdropper respectively. The proof proceeds in three steps: first it is shown that $H(\mathcal{R}|\mathcal{E}, \mathcal{U}) = 0$. The second step is to prove $H(\mathcal{E}) \leq R$. The final step shows that these two conditions suffice to guarantee complete security, i.e., $I(\mathcal{U}; \mathcal{E}) = 0$. ■

VII. MSR CODES SECURE FROM ACTIVE ADVERSARIES

MSR codes use the minimum possible storage at each node. The data-reconstruction property necessitates that the data from any k nodes suffice to reconstruct all the B message symbols. As a result, each node must necessarily store at least a fraction $\frac{1}{k}$ of the entire message. Hence for an MSR code we have $\alpha = \frac{B}{k}$. To meet the bound (1) with equality (in absence of errors/erasures), an MSR code must satisfy

$$B = k\alpha \quad (24a)$$

$$d\beta = \alpha + (k-1)\beta. \quad (24b)$$

In this section we present explicit constructions of optimal MSR codes for all parameter values $[n, k, d \geq 2k-2]$ providing on-demand security from active adversaries. We employ the striping procedure described in Section IV, and without loss of generality construct optimal codes for $\beta = 1$.

Since our codes are required to provide on-demand security, the encoding procedure is independent of the levels of security desired and is identical to the case when no security is required.

We first construct the code for the case $d = 2k - 2$ and then use the *shortening* technique of [10], [11] to obtain codes for all $d \geq 2k - 2$. As we will see subsequently, the decoding algorithms and the security properties of the code for $d > 2k - 2$ follow directly from the corresponding properties of the $d = 2k - 2$ code.

A. Encoding algorithm for $d = 2k - 2$

In the MSR regime with $d = 2k - 2$ and $\beta = 1$, from (24) we have

$$k = \alpha + 1, \quad (25a)$$

$$d = 2\alpha, \quad (25b)$$

$$B = \alpha(\alpha + 1). \quad (25c)$$

The product-matrix MSR code is described by an $(n \times \alpha)$ code matrix C of the form $C = \Psi M$, with an $(n \times d)$ encoding matrix Ψ and an $(d \times \alpha)$ message matrix M . The MSR code differs from the product-matrix MBR code in the specific design of the matrices Ψ and M .

Under the MSR regime, the $(n \times d)$ encoding matrix Ψ is of the form

$$\Psi = [\Phi \quad \Lambda\Phi], \quad (26)$$

where Φ is an $(n \times \alpha)$ matrix and Λ is an $(n \times n)$ *diagonal* matrix. The matrices Φ and Λ are chosen such that: (i) any α rows of Φ are linearly independent, (ii) any d rows of Ψ are linearly independent, and (iii) the diagonal elements of Λ are all distinct. These requirements can be met, for example, by choosing Ψ to be a Vandermonde matrix with elements chosen carefully to satisfy the third condition: the i th ($1 \leq i \leq n$) row of Ψ , $\psi_i^T = [1 \ x_i \ x_i^2 \ \cdots \ x_i^{d-1}]$, is chosen such that $x_{i_1}^\alpha \neq x_{i_2}^\alpha \ \forall 1 \leq i_1 < i_2 \leq n$. Observe that these properties make Φ the generator matrix of an $[n, \alpha]$ MDS code, and Ψ the generator matrix of an $[n, d]$ MDS code.

The choice of the matrix Ψ poses the only restriction to the choice of the finite field \mathbb{F}_q . For instance, choosing Ψ as a Vandermonde matrix (in the manner described above) permits any $q \geq n^2$.

We will now specify the design of the message matrix M . From (25b) and (25c), we have $d = 2\alpha$ and

Node	Data Stored													
Node 1	a						b							
Node 2	a	+	b	+	d	+	e	b	+	c	+	e	+	f
Node 3	a	+	3b	+	9d	+	e	b	+	3c	+	9e	+	f
Node 4	a	+	2b	+	4d	+	8e	b	+	2c	+	4e	+	8f
Node 5	a	+	6b	+	10d	+	8e	b	+	6c	+	10e	+	8f
Node 6	a	+	5b	+	12d	+	8e	b	+	5c	+	12e	+	8f
Node 7	a	+	4b	+	3d	+	12e	b	+	4c	+	3e	+	12f

Fig. 6: An MSR code for $[n = 7, k = 3, d = 4]$. The message is $\{a, b, c, d, e, f\}$ and the finite field of operation is \mathbb{F}_{13} . In the absence of security requirements ($p = 0$), repair of any node f ($1 \leq f \leq 7$) is performed by connecting to any $d = 4$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$, where x_f is the f^{th} element of $\mathbf{x} = [0 \ 1 \ 3 \ 2 \ 6 \ 5 \ 4]$. The repair of any node f ($1 \leq f \leq 7$), with security from compromise of any $p = 1$ node, is performed by connecting to any $(d + 2p) = 6$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$. The code is optimal: it requires the minimum possible storage (it is MDS), and for this storage, it requires the minimum possible amount of download in each of the aforementioned scenarios.

$B = \alpha(\alpha + 1)$ respectively. The $(d \times \alpha)$ message matrix M is constructed as

$$M = \begin{bmatrix} S_1 \\ S_2 \end{bmatrix}, \quad (27)$$

where S_1 and S_2 are $(\alpha \times \alpha)$ symmetric matrices. The matrices S_1 and S_2 together have precisely $\alpha(\alpha + 1)$ distinct entries, which are now populated by the $B = \alpha(\alpha + 1)$ message symbols. Thus, under this encoding mechanism, node i ($1 \leq i \leq n$), stores the α symbols

$$\psi_i^T M = \phi_i^T S_1 + \lambda_i \phi_i^T S_2.$$

This encoding procedure is illustrated below with an example. This example will subsequently be reused to illustrate secure node-repair.

Example 3: Let us construct an MSR code for $[n = 7, k = 3, d = 4]$. With $\beta = 1$, the associated parameters are $(\alpha = 2, B = 6)$. Note that the chosen parameters k and d satisfy the relation $d = 2k - 2$. Let us operate over the finite field \mathbb{F}_{13} , and choose the encoding matrix Ψ to be a Vandermonde matrix. Further, let us denote the 6 message symbols as $\{a, b, c, d, e, f\}$. The matrices Ψ and M take values

$$\Psi = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 3 & 9 & 1 \\ 1 & 2 & 4 & 8 \\ 1 & 6 & 10 & 8 \\ 1 & 5 & 12 & 8 \\ 1 & 4 & 3 & 12 \end{bmatrix}, \quad M = \begin{bmatrix} a & b \\ b & c \\ d & e \\ e & f \end{bmatrix}.$$

The code is given by $C = \Psi M$. The resultant code is depicted in Fig. 6.

Remark 3 (Systematic code): Although the encoding mechanism described above does not directly result in a systematic code, it can be converted easily to systematic form by following a procedure that will be described subsequently in Section VII-D.

This completes the description of the encoding procedure. We now move on to present decoding algorithms for node-repair and data-reconstruction with on-demand security.

B. Algorithm for Secure Node Repair

The following theorem presents an explicit algorithm for node-repair that ensures security in the presence of compromise of nodes to an active adversary.

Theorem 5 (Secure node-repair): In the code presented, for any desired value of p , the repair of a failed node can be secured from the compromise of up to p nodes, by letting the replacement node download $\beta = 1$ values each from $(d + 2p)$ arbitrary nodes. The amount of data downloaded in this processes is minimum, thus establishing the secrecy capacity of such systems. The choice of parameter p is arbitrary, and can be different for every instance of repair, thus providing on-demand security.

Proof: Suppose during any instance of node-repair, it is desired to provide protection from the compromise of up to p nodes, for some choice of p . Let f denote the failed node and $\psi_f = [\phi_f \ \lambda_f \phi_f]$ be its encoding vector. The α ($= d$) symbols stored in the failed node are

$$\psi_f^T M = \phi_f^T S_1 + \lambda_f \phi_f^T S_2 . \quad (28)$$

The replacement node connects to some $(d + 2p)$ nodes. The replacement node is required to recover these d symbols by downloading $\beta = 1$ symbol each from the $(d + 2p)$ nodes that it connects to, when at most p of these nodes are compromised. Let \mathcal{J} denote the set of these $(d + 2p)$ nodes. Under our repair algorithm, each of these $(d + 2p)$ nodes pass the inner-product of the α values stored in them with the encoding vector ψ_f of the failed node f . That is, for every $j \in \mathcal{J}$, node j passes the value of

$$\psi_j^T M \psi_f .$$

This value may be corrupt if node j is compromised to the adversary.

Let $\Psi_{\mathcal{J}}$ denote the $((d + 2p) \times d)$ submatrix of Ψ , with the rows of $\Psi_{\mathcal{J}}$ comprising the $(d + 2p)$ rows of Ψ that correspond to the $(d + 2p)$ nodes in \mathcal{J} . Then the replacement node has access to the $(d + 2p)$ values

$$\Psi_{\mathcal{J}} M \phi_f .$$

To simplify notation, let us define a quantity μ as

$$\mu = M \phi_f . \quad (29)$$

In terms of this notation, the $(d + 2p)$ encoded symbols downloaded by the replacement node can be written as $\Psi_{\mathcal{J}} \mu$. To decode in the presence of possible corruption, we will exploit the MDS property of the matrix Ψ . By construction, Ψ is the generator matrix of an $[n, d]$ MDS code. Hence, $\Psi_{\mathcal{J}}$ is the generator matrix of a $[d + 2p, d]$ MDS code. This implies that $\Psi_{\mathcal{J}}$ has a minimum distance of $(2p + 1)$. The $(d + 2p)$ symbols $\Psi_{\mathcal{J}} \mu$ downloaded by the replacement node are simply an encoding of μ under a code that has $\Psi_{\mathcal{J}}$ as its generator matrix. Hence, the replacement node can recover μ correctly by decoding the MDS code $\Psi_{\mathcal{J}}$ in the presence of up to p corruptions. Thus, even in the presence of up to p compromised nodes in the system, the replacement

node correctly recovers

$$\begin{aligned}\mu &= M\phi_f \\ &= \begin{bmatrix} S_1 \\ S_2 \end{bmatrix} \phi_f\end{aligned}\tag{30}$$

$$= \begin{bmatrix} S_1\phi_f \\ S_2\phi_f \end{bmatrix}\tag{31}$$

Thus the replacement node has access to the correct values of $S_1\phi_f$ and $S_2\phi_f$. The symmetry of S_1 and S_2 allows for the computation of

$$\phi_f^T S_1 + \lambda_f \phi_f^T S_2$$

correctly, and this is precisely the collection of α symbols that were stored in the failed node. ■

The following example illustrates this repair algorithm.

Example 4: Consider the code described in Example 3 (Fig. 6) for parameters $[n = 7, k = 3, d = 4]$. Suppose node 2 fails and needs to be repaired with a security from the compromise of any $p = 1$ other node. From the construction described in Example 3, we see that $\phi_f^T = [1 \ 1]$. Under the secure repair algorithm described above, the replacement node connects to the $\Delta = d + 2p = 6$ remaining nodes. Each of these nodes takes an inner product of its data with the vector $[1 \ 1]$, i.e., passes the sum of the two symbols it stores. The data thereby obtained by the replacement node is a $[6, 4]$ -MDS code with the message as $\{a+b, b+c, d+e, e+f\}$; the six symbols obtained by the replacement node are

- $a + b$ $(= (a + b) + 0(b + c) + 0(d + e) + 0(e + f))$
- $a + 3b + 9d + e + b + 3c + 9e + f$ $(= (a + b) + 3(b + c) + 9(d + e) + (e + f))$
- $a + 2b + 4d + 8e + b + 2c + 4e + 8f$ $(= (a + b) + 2(b + c) + 4(d + e) + 8(e + f))$
- $a + 6b + 10d + 8e + b + 6c + 10e + 8f$ $(= (a + b) + 6(b + c) + 10(d + e) + 8(e + f))$
- $a + 5b + 12d + 8e + b + 5c + 12e + 8f$ $(= (a + b) + 5(b + c) + 12(d + e) + 8(e + f))$
- $a + 4b + 3d + 12e + b + 4c + 3e + 12f$ $(= (a + b) + 4(b + c) + 3(d + e) + 12(e + f))$

The $[6, 4]$ MDS property allows for correction of one arbitrary error. It follows that the replacement node can correctly recover $\{a+b, b+c, d+e, e+f\}$ and compute its desired data $(a+b+d+e, b+c+e+f)$ even in the presence of an adversary who may gain control of one arbitrary node.

C. Algorithm for Secure Data Reconstruction

The following theorem presents an explicit algorithm for data-reconstruction that ensures security in the presence of nodes compromised to an active adversary.

Theorem 6 (Secure data-reconstruction): In the code presented, a data-collector can recover all the B message symbols by downloading data stored in $(k + 2p)$ arbitrary nodes, in the presence of up to p compromised nodes. The choice of parameter p is arbitrary, and can be different for every instance of data-reconstruction, thus providing on-demand security.

Proof: Consider a data-reconstruction operation that requires security from up to p compromised nodes. Then, under our protocol for data-reconstruction, a data-collector connects to $(k + 2p)$ arbitrary nodes in the system and downloads the α symbols stored in each of these nodes. Let \mathcal{I} denote this set of $(k + 2p)$ nodes.

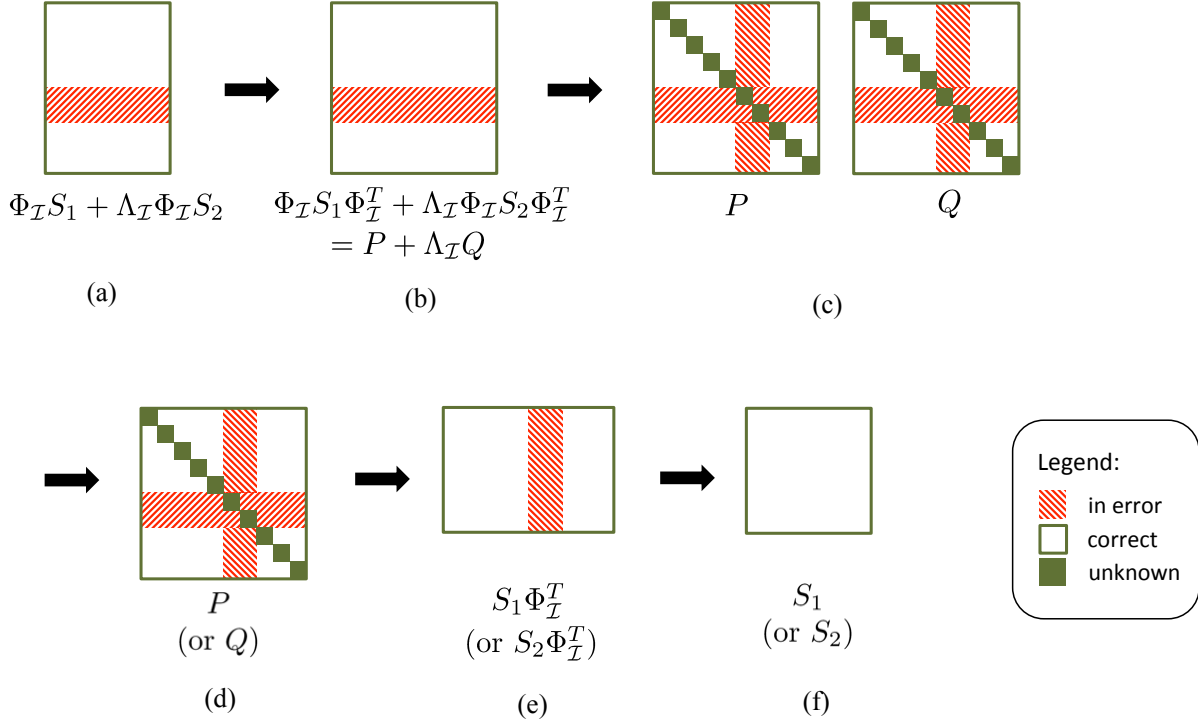


Fig. 7: An illustration of the error patterns arising during the reconstruction process under the product-matrix MSR code providing security from active adversaries. The pattern shown in part (a) depicts the corruption of some p rows (corresponding to the set \mathcal{A}) of the received data. The remaining parts (b)-(f) illustrate the resulting error patterns at various steps of the decoding algorithm. Note that the p corrupt rows are shown as contiguous only for ease of illustration.

Of these, let \mathcal{A} denote the subset of nodes compromised to the adversary. To prove the correctness of our reconstruction algorithm for the desired protection-level p , we assume that the size of \mathcal{A} is no larger than p . Of course, the set \mathcal{A} is not known to the decoding algorithm, and is used here only to illustrate the algorithm.

Let $\Psi_{\mathcal{I}}$ denote the $((k+2p) \times d)$ submatrix of Ψ , whose rows comprise the set \mathcal{I} of the rows of Ψ . From the structure of the encoding matrix Ψ (26), we can write $\Psi_{\mathcal{I}}$ as

$$\Psi_I = \begin{bmatrix} \Phi_I & \Lambda_I \Phi_I \end{bmatrix}.$$

Then the data collector has access to the encoded symbols

$$\begin{aligned}\Psi_I M &= \begin{bmatrix} \Phi_I & \Lambda_I \Phi_I \end{bmatrix} \begin{bmatrix} S_1 \\ S_2 \end{bmatrix} \\ &= \begin{bmatrix} \Phi_I S_1 + \Lambda_I \Phi_I S_2 \end{bmatrix},\end{aligned}\tag{32}$$

of which the p rows corresponding to \mathcal{A} can be corrupt. An example of such an error pattern is illustrated in Fig. 7a.

The data-collector multiplies this term by $\Phi_{\mathcal{I}}^T$ to obtain

$$\begin{aligned}\Psi_{\mathcal{I}} M \phi_{\mathcal{I}}^T &= \left[\Phi_{\mathcal{I}} S_1 + \Lambda_{\mathcal{I}} \Phi_{\mathcal{I}} S_2 \right] \Phi_{\mathcal{I}}^T \\ &= \Phi_{\mathcal{I}} S_1 \Phi_{\mathcal{I}}^T + \Lambda_{\mathcal{I}} \Phi_{\mathcal{I}} S_2 \Phi_{\mathcal{I}}^T.\end{aligned}\quad (33)$$

Note that since this operation involves only column operations, the error patterns do not change, and only the p rows corresponding to \mathcal{A} continue to be corrupt. Next, define two $((\alpha + 1) \times (\alpha + 1))$ matrices P and Q as

$$P = \Phi_{\mathcal{I}} S_1 \Phi_{\mathcal{I}}^T, \quad (34)$$

$$Q = \Phi_{\mathcal{I}} S_2 \Phi_{\mathcal{I}}^T. \quad (35)$$

Since S_1 and S_2 are symmetric, it follows that the matrices P and Q are also symmetric. One can rewrite the data available to the data-collector in terms of P and Q as

$$P + \Lambda_{\mathcal{I}} Q,$$

of which the p rows corresponding to \mathcal{A} can be corrupt (Fig. 7b). Now, for any i and j such that $1 \leq i, j \leq k+2p$ and $i \neq j$, the $(i, j)^{\text{th}}$ element of the matrix $P + \Lambda_{\mathcal{I}} Q$ is

$$P_{ij} + \lambda_i Q_{ij}, \quad (36)$$

and its $(j, i)^{\text{th}}$ element is

$$\begin{aligned}P_{ji} + \lambda_j Q_{ji} \\ = P_{ij} + \lambda_j Q_{ij},\end{aligned}\quad (37)$$

where (37) follows from the symmetry of matrices P and Q . By construction, $\lambda_i \neq \lambda_j$ whenever $i \neq j$, and hence using (36) and (37), the data-collector solves for the values of P_{ij} and Q_{ij} for all $i \neq j$ (in other words, the data-collector solves for all non-diagonal elements of P and Q). However, some of these computed values may be in error due to the corruption of p of the rows in $P + \Lambda_{\mathcal{I}} Q$. In particular, the p rows (and hence the p columns) corresponding to \mathcal{A} may be in error in P as well as in Q . These error patterns in P and Q are illustrated in Fig. 7c.

We will now use the (possibly corrupt) values of the non-diagonal elements of matrices P and Q to correctly decode S_1 and S_2 respectively. Let us first consider recovery of S_1 , for which we will use the values of the non-diagonal elements of P obtained above. Consider any column j ($1 \leq j \leq k+2p$) of the matrix P , and let $\mathcal{J} = \mathcal{I} \setminus \{j\}$. The j^{th} column of the matrix P , excluding the element on its diagonal, is

$$\Phi_{\mathcal{J}} S_1 \phi_j,$$

and is of length $(k + 2p - 1) = (\alpha + 2p)$. Recall that by construction, Φ is the generator matrix of an $[n, \alpha]$ MDS code. Hence its submatrix $\Phi_{\mathcal{J}}$ is the generator matrix of an $[\alpha + 2p, \alpha]$ MDS code. It follows that $\Phi_{\mathcal{J}}$ has a minimum distance of $(2p + 1)$. The symbols $\Phi_{\mathcal{J}} S_1 \phi_j$ can be considered as an encoding of the vector $S_1 \phi_j$ using the generator matrix $\Phi_{\mathcal{J}}$. Using the MDS property of $\Phi_{\mathcal{J}}$, the algorithm attempts to decode $S_1 \phi_j$, assuming the presence of no more than p errors. Recall from the discussion above that the vector $S_1 \phi_j$ has at most p corrupt entries if $j \notin \mathcal{A}$, but it can be entirely corrupt if $j \in \mathcal{A}$ (see Fig. 7d). As a result, the data-collector obtains the correct value of $S_1 \Phi_{\mathcal{J}}$ if $j \notin \mathcal{A}$, and a possibly corrupt value if $j \in \mathcal{A}$ as shown in Fig. 7e (of course, the data-collector still does not know \mathcal{A} and hence does not know if the value obtained is correct or not). However, when the size of \mathcal{A} is at most p , at most p of these decoded vectors $\{S_1 \phi_j\}_{j=1}^{\alpha+2p+1}$

may be in error. Now, these vectors comprise the columns of the matrix $S_1 \Phi_x^T$ as seen in Fig. 7e. Again, since Φ_x is the generator matrix of an $[\alpha + 2p + 1, \alpha]$ MDS code, one can decode the $(\alpha \times \alpha)$ matrix S_1 from $S_1 \Phi_x^T$ in the presence of up to p erroneous columns. Thus, the data-collector correctly recovers S_1 .

Following steps identical to recovery of S_1 from P , the data-collector also correctly decodes matrix S_2 from the non-diagonal elements of Q . In this manner, the data collector correctly recovers all the message symbols in the presence of at most p compromised nodes. ■

D. Conversion to systematic form

Any MSR code where failed nodes are replaced by their exact replicas can be converted to systematic form through a remapping of the source symbols. An explicit algorithm to convert the product-matrix MSR codes was provided in [11], and is briefly reproduced here for subsequent use. Let Ψ_k denote the $(k \times d)$ submatrix of matrix Ψ with its k rows comprising the first k rows of Ψ . Let U be a $(k \times \alpha)$ matrix consisting of the $B (= k\alpha)$ message symbols as its $k\alpha$ entries. Now, choose the message matrix M to satisfy

$$\Psi_k M = U, \quad (38)$$

while maintaining the structure of M as given by (27).

The value of M in (38) can be obtained by using the decoding algorithm for data reconstruction (for $p = 0$). With this M as the message matrix, under the encoding algorithm of Section VII-A, the first k nodes will store the data $\Psi_k M$. This is precisely the matrix of *uncoded* message symbols U , thus making these nodes systematic. Note that following this remapping, the node repair algorithm remains the same. Data reconstruction from the k systematic nodes requires no computation, while decoding from parity nodes requires the additional step of recovering U from M by pre-multiplication with Ψ_k as in (38).

Example 5: Fig. 8 depicts an example of a *systematic* product-matrix MSR code providing (on-demand) security from malicious adversaries. This code is obtained by converting the code of Example 3 (Fig. 6) into a systematic form by remapping the source symbols.

E. Shortening for $d \geq 2k - 2$

A shortening technique under the MSR regime for constructing low-redundancy codes from the high-redundancy codes was introduced in [10], [13]. This technique is employed in [11] to extend the $d = 2k - 2$ product-matrix MSR code to the parameter range of $d \geq 2k - 2$. In this section we construct secure MSR codes for $d \geq 2k - 2$ that meet (24) by applying this shortening technique to the secure MSR codes of Section VII-A.

Since the codes are required to provide on-demand security, the shortening procedure to obtain secure codes for $d \geq 2k - 2$ is independent of the desired security levels, and is *identical* to the corresponding procedure in the absence of security [11, Section V-C]. We briefly describe this procedure and show that these codes provide optimal security from active adversaries.

Suppose for some positive integer i , one wishes to construct a product-matrix MSR code with parameters $[n, k, d = 2k - 2 + i]$. Moreover, it is desired to construct an optimal code, i.e., a code whose parameters (B, α, β) satisfy (24). As the first step in this task, we construct a product-matrix MSR code for the parameters $[n' = n + i, k' = k + i, d' = d + i = 2k - 2 + 2i]$. Note that this parameter set satisfies $d' = 2k' - 2$, and one can employ the encoding procedure described in Section VII-A for this construction. Next, this code is converted to a systematic form as described in Section VII-D. Let us denote the resultant systematic code by C' . In order

Node	Data Stored																	
Node 1	a						b											
Node 2	c						d											
Node 3	e						f											
Node 4	6a	+	8b	+	2c	+	6e	6a	+	4b	+	3c	+	11d	+	4e	+	10f
Node 5	6a	+	4b	+	11c	+	10e	3a	+	5b	+	8c	+	9d	+	2e	+	12f
Node 6	6a	+	9b	+	c	+	7e	10a	+	7b	+	5c	+	3d	+	11e	+	5f
Node 7	2a	+	4b	+	5c	+	7e	3a	+	b	+	8c	+	3d	+	2e	+	9f

Fig. 8: A *systematic* MSR code for $(n = 7, k = 3, d = 4)$. The message is $\{a, b, c, d, e, f\}$ and the finite field of operation is \mathbb{F}_{13} . In the absence of security requirements ($p = 0$), repair of any node f ($1 \leq f \leq 7$) is performed by connecting to any $d = 4$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$, where x_f is the f^{th} element of $\mathbf{x} = [0 \ 1 \ 3 \ 2 \ 6 \ 5 \ 4]$. The repair of any node f ($1 \leq f \leq 7$), with security from compromise of any $p = 1$ node, is performed by connecting to any $(d + 2p) = 6$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$. The code is optimal: it requires the minimum possible storage (it is MDS), and for this storage, it requires the minimum possible amount of download in each of the aforementioned scenarios.

to obtain the desired $[n, k, d = 2k - 2 + i]$ code \mathcal{C} , we observe from (24) that the number of message symbols B' in code \mathcal{C}' is $(k + i)(k - 1 + i)$, and the number of message symbols B in code \mathcal{C} is $k(k - 1 + i)$. Taking a cue from this expression, to obtain code \mathcal{C} from code \mathcal{C}' , we set the $i(k - 1 + i)$ message symbols stored in the first i (systematic) nodes in code \mathcal{C}' to zero. The data stored in the remaining n nodes, which is an encoding of the remaining $B = k(k - 1 + i)$ message symbols, constitutes the desired $[n, k, d = 2k - 2 + i]$ code \mathcal{C}' .

Example 6: Fig. 9 depicts an example of a product-matrix MSR code providing (on-demand) security from malicious adversaries for a parameter set satisfying $d > 2k - 2$. This code is obtained by shortening the code of Example 5 (Fig. 8) by eliminating message symbols $\{a, b\}$ from the code of Example 5. The message symbols $\{c, d, e, f\}$ in the code of Example 5 (Fig. 6) correspond to $\{a, b, c, d\}$ respectively in Fig. 9.

The following theorem provides explicit algorithms for secure reconstruction and repair algorithms for this code.

Theorem 7 (Secure data-reconstruction and node-repair): In the code presented, a data-collector can recover all the B message symbols by downloading data stored in $(k + 2p)$ arbitrary nodes, and a replacement node can recover all the data stored in the failed node by downloading $\beta = 1$ symbols each from $(d + 2p)$ arbitrary nodes, in the presence of up to p compromised nodes. The choice of parameter p is arbitrary, and can be different for every instance of data-reconstruction and node-repair, thus providing on-demand security.

Proof: To see how reconstruction and repair are performed in code \mathcal{C} , we can pretend to operate under code \mathcal{C}' , and assume that a user (or replacement node) always connects to the first i nodes in addition to the k (or d) nodes that it chooses in \mathcal{C} . The data in the first i nodes is known to be all zero, and hence cannot be erroneous. Thus, reconstruction or repair in \mathcal{C} with a security level of p is identical to that in \mathcal{C}' with the same security level p . It follows that the reconstruction and repair operations in code \mathcal{C} are identical to those in \mathcal{C}' .

Optimality is a result of the parameters of \mathcal{C} satisfying the bound (24). ■

Node	Data Stored			
Node 1	a		b	
Node 2	c		d	
Node 3	2a	+	6c	3a + 11b + 4c + 10d
Node 4	11a	+	10c	8a + 9b + 2c + 12d
Node 5	a	+	7c	5a + 3b + 11c + 5d
Node 6	5a	+	7c	8a + 3b + 2c + 9d

Fig. 9: A systematic MSR code for $(n = 6, k = 2, d = 3)$. The message is $\{a, b, c, d\}$ and the finite field of operation is \mathbb{F}_{13} . In the absence of security requirements ($p = 0$), repair of any node f ($1 \leq f \leq 6$) is performed by connecting to any $d = 3$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$, where x_f is the f^{th} element of $\mathbf{x} = [0 \ 1 \ 2 \ 6 \ 5 \ 4]$. The repair of any node f ($1 \leq f \leq 6$), with security from compromise of any $p = 1$ node, is performed by connecting to any $(d + 2p) = 5$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$. The code is optimal: it requires the minimum possible storage (it is MDS), and for this storage, it requires the minimum possible amount of download in each of the aforementioned scenarios.

VIII. MSR CODES FOR SECURITY FROM PASSIVE EAVESDROPPERS

Recall from (24) that in the absence of secrecy requirements, the MSR regime has

$$d\beta = \alpha + (k - 1)\beta. \quad (39)$$

Unlike the MSR codes for security from active adversaries, the codes for security from passive eavesdroppers do not provide on-demand security. Thus, one needs to fix the desired level of security $\{\ell, m\}$ at the time of encoding.

In this section we present explicit constructions of MSR codes for all parameter values $[n, k, d \geq 2k - 2]$ and all $\{\ell, m\}$ providing information-theoretic security from passive eavesdroppers. The $\{\ell, m\}$ -secure MSR codes constructed in this paper achieve

$$\begin{aligned} B^* &= (k - \ell)(\alpha - m\beta) \\ &= (k - \ell)(d - k + 1 - m)\beta. \end{aligned} \quad (40)$$

Note that (40) meets the lower bound (7) for $m \leq 1$ making our codes optimal for this regime and also establishing the secure capacity of this parameter regime.

For a specified set of parameters $[n, k, d]$, let R denote the difference in the number of message symbols B that can be stored in a system without security, and the number B^* that can be stored in the $\{\ell, m\}$ -secure system constructed here. From (24) and (40), the value of R is given by

$$\begin{aligned} R &= B - B^* \\ &= m(k - \ell)\beta + \ell\alpha. \end{aligned} \quad (41)$$

We now describe the construction of the secure product-matrix MSR code, which is performed by replacing precisely R message symbols in the code of VII-A by R random symbols. As done in Section VII, we will first consider the case of $d = 2k - 2$, and subsequently extend it to $d \geq 2k - 2$ via a shortening procedure.

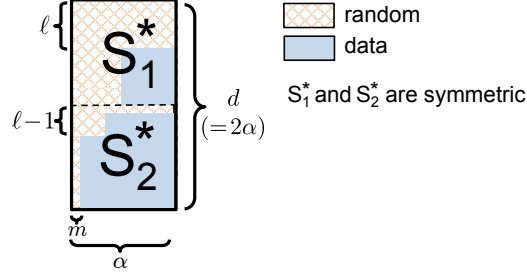


Fig. 10: Structure of the message matrix M^* for an MSR code secure from passive eavesdroppers, when $d = 2k - 2$. The matrix M^* for the code providing security from passive eavesdroppers is obtained from the message matrix M of the code without security requirements by filling up a specific set of positions (which were occupied by message values in M) by random values. The two submatrices S_1^* and S_2^* are symmetric.

Observe from (39) and (40) that the parameters B^* and α are multiples of β . This allows us to perform striping (explained in Section IV) and construct codes for $\beta = 1$ without loss of generality.

A. Encoding for $d = 2k - 2$

An $\{\ell, m\}$ -secure product-matrix MSR code for the parameters $[n, k, d]$ is obtained by modifying the $[n, k, d]$ product-matrix MSR code constructed in Section VII-A. Let us denote the product-matrix MSR code of Section VII-A by \mathcal{C} , and the code with security (which will be constructed below) by \mathcal{C}^* .

The code \mathcal{C}^* also belongs to the product-matrix framework, and is described by an $(n \times \alpha)$ code matrix $C^* = \Psi^* M^*$. The matrices Ψ^* and M^* are obtained by modifying the matrices Ψ and M of Section VII-A as follows. Choose Ψ^* such that it satisfies the following property in addition to those required for Ψ : when restricted to the first ℓ columns, any ℓ rows of Ψ^* are linearly independent. The choice of Ψ^* as a Vandermode matrix as described in Section VII-A satisfies this additional property.

When $\beta = 1$, the value of R in (41) equals $R = \ell\alpha + (k - \ell)m$. The message matrix M^* for code \mathcal{C}^* is obtained by replacing a specific set of R message symbols in M with random symbols. Recall the structure of matrix M^* from (27). Matrix M^* is obtained from M by replacing: the $\left(\ell\alpha - \frac{\ell(\ell-1)}{2}\right)$ symbols in the first ℓ rows (and hence the first ℓ columns) of the symmetric matrix S_1 , the $\frac{\ell(\ell-1)}{2}$ symbols in the *intersection* of the first $(\ell - 1)$ rows and first $(\ell - 1)$ columns of the symmetric matrix S_2 , and the $(k - \ell)m$ remaining symbols in the first m rows (and hence the first m columns) of S_2 . Fig. 10 depicts this procedure.

The following example illustrates this encoding procedure.

Example 7: In this example, we will construct an MSR code for $[n = 6, k = 3, d = 4]$. Let the alphabet of operation be \mathbb{F}_{13} . With $\beta = 1$, (24) leads to having $\alpha = 2$ and $B = 6$ in the absence of security requirements. Let us suppose we wish to provide security from an eavesdropper who can read the data stored in any one arbitrary node but not read any data downloaded during repair operations. This corresponds to parameters $\ell = 1$ and $m = 0$. The maximum size of the message that can be stored (40) is $B^* = 4$. We thus choose $B - B^* = 2$ symbols r_1 and r_2 uniformly at random from \mathbb{F}_{13} . Let us denote the 4 message symbols as $\{a, b, c, d\}$. In

Node	Data Stored													
Node 1	r_1						r_2							
Node 2	r_1	+	r_2	+	b	+	c	r_2	+	a	+	c	+	d
Node 3	r_1	+	$3r_2$	+	$9b$	+	c	r_2	+	$3a$	+	$9c$	+	d
Node 4	r_1	+	$2r_2$	+	$4b$	+	$8c$	r_2	+	$2a$	+	$4c$	+	$8d$
Node 5	r_1	+	$6r_2$	+	$10b$	+	$8c$	r_2	+	$6a$	+	$10c$	+	$8d$
Node 6	r_1	+	$5r_2$	+	$12b$	+	$8c$	r_2	+	$5a$	+	$12c$	+	$8d$
Node 7	r_1	+	$4r_2$	+	$3b$	+	$12c$	r_2	+	$4a$	+	$3c$	+	$12d$

Fig. 11: An MSR code for $(n = 7, k = 3, d = 4)$ providing $\{\ell = 1, m = 0\}$ security from passive eavesdroppers: the data is secure from passive eavesdroppers who can read the data *stored* in any one arbitrary node but not to the data passed during repair of any node. The message is $\{a, b, c\}$ and the finite field of operation is \mathbb{F}_{13} . The symbols r_1 and r_2 are drawn uniformly (and independently) at random from \mathbb{F}_{13} . The data of any node f ($1 \leq f \leq 6$), when required to protect from compromise of any $p = 1$ node, is obtained by connecting to any $(d + 2p) = 6$ nodes and downloading the inner product of the two symbols in each node with $[1 \ x_f]$ where x_f is the f^{th} element of $\mathbf{x} = [0 \ 1 \ 3 \ 2 \ 6 \ 5 \ 4]$. The code is optimal in the sense that the amount of storage required is the minimum possible in this setting (it is MDS), and furthermore, the amount of download required for any repair is also the minimum possible under this amount of storage.

this code, we choose the matrices Ψ^* , S_1^* , S_2^* and M^* as

$$\Psi^* = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 3 & 9 & 1 \\ 1 & 2 & 4 & 8 \\ 1 & 6 & 10 & 8 \\ 1 & 5 & 12 & 8 \\ 1 & 4 & 3 & 12 \end{bmatrix}, \quad S_1^* = \begin{bmatrix} r_1 & r_2 \\ r_2 & a \end{bmatrix}, \quad S_2^* = \begin{bmatrix} b & c \\ c & d \end{bmatrix}, \quad M^* = \begin{bmatrix} S_1 \\ S_2 \end{bmatrix}.$$

The code is given by $C^* = \Psi^* M^*$. The data stored by the seven storage nodes under this code is depicted in Fig. 11. One can verify that under this code, the data stored in any single node provides no information about the message.

B. Reconstruction, Repair and Security

The following theorems prove the properties of reconstruction, repair and security in the secure PM-MSR code.

Theorem 8 (Data-reconstruction and node-repair): In the code presented, a data-collector can recover all the B^* message symbols by downloading the data stored in any k nodes, and a failed node can be repaired by downloading $\beta = 1$ symbol each from any d remaining nodes.

Proof: As in the proof of Theorem 3, treating the random symbols also as message symbols, the secure PM-MSR code C^* becomes identical to the PM-MSR code C . Thus reconstruction and repair in C^* are identical to that in C . \blacksquare

Theorem 9 (Information-theoretic security): The code C^* ensures $\{\ell, m\}$ -security, i.e., an eavesdropper hav-

ing access to the data stored in up to ℓ nodes, and to all the data passed to up to m of these nodes during one or more of their repair operations, gets no information about the message.

Proof: The proof follows the three-step procedure as in the proof in the MBR case (Theorem 4). Please see the Appendix for the complete proof. ■

C. Shortening for $d \geq 2k - 2$

For any $i \geq 1$, we will now construct product-matrix MSR codes having the parameters $[n, k, d = 2k - 2 + i]$, for any desired level $\{\ell, m\}$ of security. Let us denote this code (which we will construct below) as \mathcal{C}^* . The code \mathcal{C}^* will satisfy (40), and is optimal when $m \in \{0, 1\}$. Throughout, we will assume without loss of generality that the data is striped, and consider $\beta = 1$ without loss of generality.

As the first step in this task, a secure product-matrix MSR code is constructed for the parameters $[n' = n + i, k' = k + i, d' = d + i = 2k - 2 + 2i]$, with a security level $\{\ell + i, m\}$. Note that this parameter set satisfies $d' = 2k' - 2$, and one can employ the encoding procedure described in Section VIII-A for this construction. Denote this code as \mathcal{C}'^* .

Our goal is to derive the desired code \mathcal{C}^* by shortening the code \mathcal{C}'^* . The shortening procedure described in Section VII-E provides one candidate approach towards this goal. Under this procedure, the code \mathcal{C}^* is derived from \mathcal{C}'^* by setting the data in the first i nodes to zero, and treating all operations in the resultant code \mathcal{C}^* as operations in \mathcal{C}'^* . However, this procedure requires the code to be made systematic, which is not possible under a code that provides security, since the compromise of any of the systematic nodes would directly reveal information about the message. Thus, in this section, we will follow a slightly different route to achieve that goal.

Let matrices Ψ'^* and M'^* respectively denote the $(n' \times d')$ encoding matrix and the $(d' \times d')$ message matrix of code \mathcal{C}'^* . We will require the matrix Ψ'^* to satisfy an additional property that when restricted to the first i columns, any set of i rows are linearly independent. The choice of Ψ'^* as a Vandermonde matrix, as done in Section VII-A, satisfies this condition.

Now, for the specific parameters of codes \mathcal{C}^* and \mathcal{C}'^* , observe from (40) that

$$B^* = B'^* = (k - \ell)(k - 1 + i - m) .$$

However, from (41), we see that the number of random symbols R^* in code \mathcal{C}^* is

$$R^* = \ell(k - 1 + i) + (k - \ell)m$$

and the number of random symbols R'^* in \mathcal{C}'^* is

$$\begin{aligned} R'^* &= (\ell + i)(k - 1 + i) + (k - \ell)m \\ &= i(k - 1 + i) + R^* \\ &= i\alpha + R^* . \end{aligned}$$

Here, $\alpha = (k - 1 + i) = (d - k + 1) = (d' - k' + 1)$ is the storage capacity per node in both codes \mathcal{C}^* and \mathcal{C}'^* .

To obtain code \mathcal{C}^* from code \mathcal{C}'^* , we replace a subset of $i\alpha$ random symbols from M'^* by a set of deterministically chosen symbols in the following manner. Let $\Psi_i'^*$ denote the $(i \times d)$ submatrix of Ψ'^* with its rows comprising the first i rows of Ψ'^* . Consider the following $i\alpha$ entries in M'^* : (a) the $i\alpha - \binom{i}{2}$ symbols in the first i rows (and columns) of the symmetric matrix $S_1'^*$, and (b) the $\binom{i}{2}$ symbols in the *intersection* of

the first $(i - 1)$ rows and the first $(i - 1)$ columns of $S_2'^*$. In the construction of code C'^* as per the procedure of Section VIII-A, each of these entries is populated by a randomly chosen value. With the goal of setting the data in first i nodes as zero, we instead choose these $i\alpha$ elements in a deterministic manner: once all other values of M'^* are fixed, we choose these values to satisfy

$$\Psi_i'^* \begin{bmatrix} S_1'^* \\ S_2'^* \end{bmatrix} = 0 .$$

This is feasible because, by construction, the first i columns of $\Psi_i'^*$ are linearly independent. Let M^* denote the $(d' \times d')$ matrix resulting from this choice.

The encoding matrix Ψ^* of code C^* is chosen to be an $(n \times d')$ matrix comprising the last n rows of Ψ'^* . Finally, the $[n, k, d = 2k - 2 + i]$ code C^* with security $\{\ell, m\}$ is given by $C^* = \Psi^* M^*$.

Theorem 10 (Data-reconstruction and node-repair): In the code presented, a data-collector can recover all the message symbols by downloading the data stored in any k nodes, and a failed node can be repaired by downloading $\beta = 1$ symbol each from any d remaining nodes.

Proof: To see how reconstruction and repair are performed in code C^* , we can pretend to operate under code C'^* , and assume that a user (or replacement node) always connects to the first i nodes in addition to the k (or d) nodes that it chooses in C^* . This is a valid assumption since the data in the first i nodes of C'^* is known to be all zero. Thus, reconstruction or repair in C^* is identical to that in C'^* , and is successful since C'^* is simply an $[n' = n + i, k' = k + i, d' = 2k' - 2 = d + i]$ product-matrix MSR code. ■

Theorem 11 (Information-theoretic security): In the code presented, an eavesdropper having access to the data stored in up to ℓ nodes, and to all the data passed to up to m of these nodes during one or more of their repair operations, gets no information about the message.

Proof: Please see the Appendix. ■

IX. NECESSARY AND SUFFICIENT CONDITIONS TO SECURE ANY REGENERATING CODE FROM ACTIVE ADVERSARIES

The conversion of the product-matrix codes into codes providing on-demand security, as described in Section V and Section VII, raises a natural question as to whether *any* regenerating code can provide on-demand security. We answer this question by providing a necessary and sufficient condition for the same. To the best of our knowledge, of all codes known to date, the only codes that satisfy this condition are the product-matrix codes [11].

Theorem 12: An $[n, k, d]$ regenerating code satisfying (2) can provide on-demand security from active adversaries and satisfy (9) with equality if and only if it has a repair mechanism that obeys the following condition: during any instance of repair, the data passed by an existing node to the replacement node does not depend on the identities of the other $(d - 1)$ nodes helping in this repair.

Proof: Necessity: Consider an $[n, k, d]$ regenerating code, and assume for now that $(n - d)$ is an odd number. Consider repair of a failed node f . Since the code provides on-demand security, it must be able to offer security from the compromise of $p = \frac{n-d-1}{2}$ nodes, by allowing the replacement node to download data from the remaining $(d + 2p) = (n - 1)$ nodes. First, observe that all the remaining $(n - 1)$ nodes help in the repair, and this removes the dependence on the answer to “which other nodes help in the repair”. Let $\sigma_1, \dots, \sigma_{n-1}$ be the data downloaded from the $(n - 1)$ remaining nodes. Since arbitrary errors in any p of these must be correctable, it follows that in the absence of any errors, any d -sized subset of $\{\sigma_1, \dots, \sigma_{n-1}\}$ must

suffice for correct decoding. Now consider the operation of this code in the absence of security requirements. In this setting, consider the repair scheme wherein a node helping in the repair of a failed node passes precisely what it passed in the case with security. The arguments above ensure that the repair will be successful, and the data σ_i passed is independent of the identities of the other $(d-1)$ nodes helping in this repair. If the parameters of the secure code satisfy (9) with equality, then those of the code without security will satisfy (2).

Recall the error detection requirement of on-demand security discussed in Remark 1. When $(n-d)$ is even, we invoke this requirement of detection of errors from the compromise of any $p = n - d - 1$ nodes, wherein the replacement node downloads data from the remaining $(d+p) = (n-1)$ nodes. The rest of the argument is identical to the case of $(n-d)$ being odd considered above.

Sufficiency: Consider a regenerating code satisfying the property that in the absence of security requirements, for repair of any node f , the data passed by any helper node h is independent of the identities of the remaining $(d-1)$ helper nodes. For a desired security level p during repair of node f , let the replacement node connect to some $(d+2p)$ arbitrary nodes. Let each of these $(d+2p)$ nodes pass the data it would have passed to the replacement node in the absence of security requirements. Since the data passed by any d of these nodes would suffice to recover the desired contents in the absence of errors, it follows that the aggregate data from the $(d+2p)$ nodes can correct p arbitrary errors. As a result, this data allows for recovery of the desired data correctly even in the presence of up to p compromise nodes. An identical argument holds for the reconstruction property, thus ensuring that the code provides on-demand security. If the parameters of the original regenerating code satisfy (2) then those of the secure code will satisfy (9) with equality. ■

Corollary 13: A code satisfying conditions in Theorem 12 must necessarily operate in the parameter regime $d < (n-1)$.

Proof: Clearly, since the number of nodes contacted during repair must satisfy $(d+2p) \leq (n-1)$, the requirement of supporting $p > 0$ requires that $n > (d+1)$. It follows that the code should not restrict the number of nodes n to be $(d-1)$. ■

Remark 4: The only explicit regenerating codes that support $n > (d+1)$ are the high-rate ‘approximately-exact’ MSR codes of [9] and the product-matrix codes [11]. However, the MSR codes of [9] do not satisfy the condition provided in Theorem 12, while as shown in Section V and Section VII, the product-matrix codes do satisfy this condition.

X. DISCUSSION

The many recent incidents of the compromise of storage systems (e.g., [29]–[31]) underscore the importance of securing the data in distributed storage systems. Such systems presently employ cryptographic techniques for providing security from passive eavesdroppers. Cryptographic techniques generally assume the eavesdropper to possess only a bounded computational power. However in practice, upon gaining access to the encrypted data, adversaries employ techniques such as intelligent guesswork, dictionary attacks or even crowd-sourcing, and often succeed in decoding significant parts of the data. Information-theoretic security, on the other hand, allows the adversary to possess unbounded computational power, but assumes that the adversary can gain access to only a limited amount of data. In the setting of distributed storage, the information-theoretic security relies on the assumption that the adversary cannot gain access to data stored in more than a certain number of nodes. However, information-theoretic security typically necessitates much greater resource-overheads; in our setting, these are overheads in storage and bandwidth.

The cryptographic and information-theoretic approaches for security rely on different kinds of assumptions,

and the security in distributed storage systems may be enhanced by employing both approaches in tandem. To see this, consider a system in which the data is first encoded using conventional cryptographic techniques, following which, an information-theoretically secure erasure code is employed for storing the data and handling node-failures. In such a setting, the adversary first needs to gain access to a certain minimum number of nodes (due to the information-theoretic security), without which it can neither obtain nor corrupt any data; even upon managing to gain such an access, the adversary obtains only the encrypted data. As a result, the layer of information-theoretic security raises the barrier to gain access to the encrypted data. Furthermore, the security provided by the cryptographic encoding allows for the use of only a *thin* layer of information-theoretic security, wherein the threshold on the number of nodes (ℓ , m or p in our setting) can be as low as 2 or 3. This thus aids in overcoming the barrier of the significant overheads faced by information-theoretic security.

An analogy from an end user's perspective is that of storing data (securely) in cloud-based storage services. For instance, one may encode encrypted data using an information-theoretically secure code, and store fractions of it in Dropbox, Google Drive and Microsoft Skydrive storage services.² Such a code will ensure that a compromise of one of these services provides zero information to the adversary. To get any information about the data, the adversary will need to hack into at least two of the three services.

In this paper, we addressed the information theoretic part. We presented explicit codes for distributed storage that offer protection from active adversaries and passive eavesdroppers. For a large range of parameters the codes achieve the outer bounds on the system requirements, thus also establishing the capacity of such systems.

ACKNOWLEDGEMENTS

The research of Nihar B. Shah was supported in part by a Berkeley Fellowship and in part by a Microsoft Research PhD fellowship. The research of K. V. Rashmi was supported in part by a Facebook Fellowship and in part by a Microsoft Research PhD Fellowship. This work was also supported in part by NSF grant CIF-0964018 and in part by NSF grant CIF-1116404. The research of P. Vijay Kumar was supported in part by NSF Grant CCF 1421848 and in part by the joint UGC-ISF Research Grant No. 1676/14.

The authors would like to thank Salim El Rouayheb and Sameer Pawar for fruitful discussions.

APPENDIX

We now present proofs of security from passive eavesdroppers. We use $H(\cdot)$ to denote Shannon entropy and $I(\cdot; \cdot)$ to denote mutual information. All logarithms are taken to base q where q is the size of the finite field of operation of the code under consideration.

Proof of Theorem 4 ($\{\ell, m\}$ security under the MBR code): We consider the worst case where the eavesdropper indeed gains access to the data stored in some ℓ nodes and the data passed during any of the repair operations of some m of these ℓ nodes.

Consider the repair of some node f under the product-matrix MBR code, and consider any node i helping in the repair process. Let ψ_f^* and ψ_i^* denote the encoding vectors of nodes f and i respectively. The data passed by node i under the repair algorithm is $(\psi_i^*)^T M^* \psi_f^*$. Since matrix M^* is symmetric, this is equal to $(\psi_f^*)^T M^* \psi_i^*$. Thus, an eavesdropper listening to the data downloaded for repair of node f obtains no more than $(\psi_f^*)^T M^*$, which is the data stored in node f .

²dropbox.com, drive.google.com, skydrive.live.com.

Let Ψ_{eve}^* be the $(\ell \times d)$ submatrix of Ψ^* , with its ℓ rows comprising the ℓ rows of Ψ corresponding to the nodes to which the eavesdropper has gained access. Thus the eavesdropper has access to the ℓd symbols in the $(\ell \times d)$ matrix E^* defined as

$$E^* = \Psi_{\text{eve}}^* M^* . \quad (42)$$

Let \mathcal{E} denote the set of these ℓd symbols that the eavesdropper has gained access to. Further, let \mathcal{U} denote the set of all B^* message symbols and let \mathcal{R} denote the set of all $R = B - B^*$ random symbols introduced at the encoding stage. In the proof, with some abuse of notation, we will also use the three terms \mathcal{E} , \mathcal{U} and \mathcal{R} to denote the random variables corresponding to the respective sets.

As the first step of the proof, we will now show that given the message symbols as side information, an eavesdropper can decode all the random symbols. Under the temporary assumption that the eavesdropper has access to all message symbols, the linearity of the code as in (42) allows the eavesdropper to subtract the effect of all message symbols from E^* . Denote the resultant $(\ell \times d)$ matrix as \tilde{E}^* . Define \tilde{M}^* as a $(d \times d)$ matrix obtained by setting all message symbols in M^* to zero. Thus \tilde{M}^* has its first ℓ rows and first ℓ columns identical to that of M^* , and zeros elsewhere. Then, given the message symbols, the eavesdropper now has access to

$$\tilde{E}^* = \Psi_{\text{eve}}^* \tilde{M}^* . \quad (43)$$

Recall the property of Ψ_{eve}^* that any ℓ rows, when restricted to the first ℓ columns, are independent. Thus, recovering the R random symbols from \tilde{E}^* is identical to data reconstruction in a product-matrix MBR code designed for $[\tilde{n} = n, \tilde{k} = \ell, \tilde{d} = d]$ in the absence of security requirements. As a result, given the message symbols, the eavesdropper can decode all the random symbols. Hence we have

$$H(\mathcal{R}|\mathcal{E}, \mathcal{U}) = 0 . \quad (44)$$

We will now show that $H(\mathcal{E}) \leq R$. From the value of R in (23), it suffices to show that out of the ℓd symbols that the eavesdropper has access to, $\frac{\ell(\ell-1)}{2}$ of them are functions (linear combinations) of the rest. To this end, consider the $(\ell \times \ell)$ matrix

$$E^*(\Psi_{\text{eve}}^*)^T = \Psi_{\text{eve}}^* M^*(\Psi_{\text{eve}}^*)^T . \quad (45)$$

Since M^* is symmetric, the $(\ell \times \ell)$ matrix in (45) is also symmetric. Thus $\binom{\ell}{2}$ dependencies among the elements of E^* can be described by the $\binom{\ell}{2}$ upper-triangular elements of the expression

$$E^*(\Psi_{\text{eve}}^*)^T - \Psi_{\text{eve}}^*(E^*)^T = 0 . \quad (46)$$

The linear-independence of the rows of Ψ_{eve}^* implies that these $\frac{\ell(\ell-1)}{2}$ (redundant) equations are linearly independent. Thus the eavesdropper has access to at most $\left(\ell d - \frac{\ell(\ell-1)}{2}\right)$ independent symbols, i.e.,

$$H(\mathcal{E}) \leq R . \quad (47)$$

The final part of the proof establishes that given (44) and (47), the eavesdropper obtains no information about

the message.

$$I(\mathcal{U}; \mathcal{E}) = H(\mathcal{E}) - H(\mathcal{E}|\mathcal{U}) \quad (48)$$

$$\leq R - H(\mathcal{E}|\mathcal{U}) \quad (49)$$

$$= R - H(\mathcal{E}|\mathcal{U}) + H(\mathcal{E}|\mathcal{U}, \mathcal{R}) \quad (50)$$

$$= R - I(\mathcal{E}; \mathcal{R}|\mathcal{U}) \quad (51)$$

$$= R - (H(\mathcal{R}|\mathcal{U}) - H(\mathcal{R}|\mathcal{E}, \mathcal{U})) \quad (52)$$

$$= R - H(\mathcal{R}|\mathcal{U}) \quad (53)$$

$$= R - H(\mathcal{R}) \quad (54)$$

$$= R - R \quad (55)$$

$$= 0, \quad (56)$$

where (49) follows from (47); (50) is a result of the fact that every symbol in the system is a function of \mathcal{U} and \mathcal{R} , resulting in $H(\mathcal{E}|\mathcal{U}, \mathcal{R}) = 0$; (53) follows from (44); and (54) follows from the fact that the random symbols are independent of the message symbols. ■

Proof of Theorem 9 ($\{\ell, m\}$ security under the MSR code when $d = 2k - 2$): We consider the worst case where the eavesdropper indeed gains access to the data stored in some ℓ nodes and the data passed during any of the repair operations of some m of these ℓ nodes. Let Ψ_ℓ^* be the $(\ell \times d)$ submatrix of Ψ^* , with its ℓ rows comprising the ℓ rows of Ψ corresponding to the nodes to which the eavesdropper has gained access. Due to the specific structure of Ψ^* as $\Psi^* = [\Phi^* \quad \Lambda^* \Phi^*]$, one can write Ψ_ℓ^* as $\Psi_\ell^* = [\Phi_\ell^* \quad \Lambda_\ell^* \Phi_\ell^*]$ where Φ_ℓ^* and Λ_ℓ^* are corresponding submatrices of Φ^* and Λ^* , and are sized $(\ell \times \alpha)$ and $(\ell \times \ell)$ respectively. Further, let Φ_m^* be the $(m \times \alpha)$ submatrix of Φ^* , corresponding to the m nodes whose repair operations are also eavesdropped upon. These m nodes are a subset of the set of ℓ nodes that constitute the matrix Ψ_{eve}^* , and hence, Φ_m^* is a submatrix of Φ_ℓ^* .

Let \mathcal{E} denote the set of symbols that the eavesdropper has gained access to. Further, let \mathcal{U} denote the set of all B^* message symbols and let \mathcal{R} denote the set of all $R = B - B^*$ random symbols introduced at the encoding stage. In the proof, with some abuse of notation, we will also use the three terms \mathcal{E} , \mathcal{U} and \mathcal{R} to denote the random variables corresponding to the respective sets.

Under the repair algorithm of the product-matrix MSR code (see Theorem 8), \mathcal{E} comprises the elements of the $(\ell \times \alpha)$ matrix $\Psi_\ell^* M^*$ and the elements of the $(d \times m)$ matrix $M^* (\Phi_m^*)^T$.

Following the approach described in Section II-B, we will first show that given the message symbols as side information, an eavesdropper can decode all the random symbols. Next, using the properties of the matrix Ψ^* and the specific structure of the message matrix M^* , we will show that $H(\mathcal{E}) \leq R$. At this point, the arguments in (48) to (56) establish that the eavesdropper obtains no information about the message.

Step 1: In this step, we show that $H(\mathcal{R}|\mathcal{E}, \mathcal{U}) = 0$. To this end, assume that a genie provides the values of all message symbols \mathcal{U} to the eavesdropper. Since the code is linear, the eavesdropper can subtract the effects of the message symbols from each of the symbols in \mathcal{E} . This allows us to assume (for this step) the values of all message symbols to be zero. Define matrix

$$\tilde{M}^* = \begin{bmatrix} \tilde{S}_1^* \\ \tilde{S}_2^* \end{bmatrix}$$

as being identical to M^* , but with all message symbols set to zero. It suffices to show that under a code operating

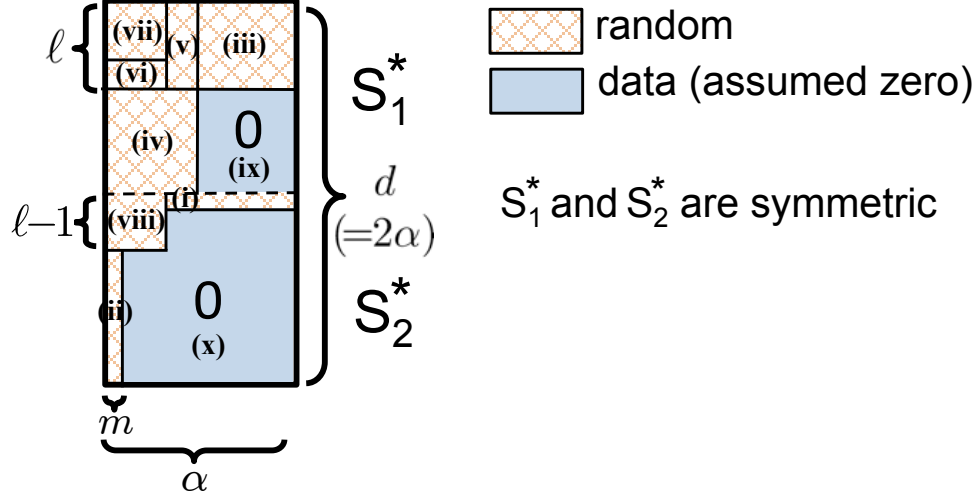


Fig. 12: Illustration of the order of decoding of the different parts of the message matrix M^* in Step 1 of the proof of security for MSR codes. Note that the parameters considered here satisfy $d = 2k - 2 = 2\alpha$.

with \tilde{M}^* as its message matrix, the eavesdropper can decode all the R random symbols. The eavesdropper has access to the data $\Psi_\ell^* \tilde{M}^*$ and $\tilde{M}^* (\Phi_m^*)^T$, and hence equivalently to:

$$\tilde{E}_1^* = \Phi_m^* \tilde{S}_1^*$$

$$\tilde{E}_2^* = \Phi_m^* \tilde{S}_2^*$$

$$\tilde{E}_3^* = \Phi_\ell^* \tilde{S}_1^* + \Lambda_\ell \Phi_\ell^* \tilde{S}_2^*$$

where \tilde{E}_1^* and \tilde{E}_2^* are both sized $(m \times (k - 1))$ and \tilde{E}_3^* is sized $(\ell \times (k - 1))$. By construction, the intersection of the rightmost $(k - \ell)$ columns and the bottommost $(k - 1 - m)$ rows of the $((k - 1) \times (k - 1))$ matrix \tilde{S}_2^* contains only the message symbols (depicted as (x) in Fig. 12) which are set to zero under our assumptions. Since the rows of Φ_m^* are linearly independent, one can decode the rightmost $(k - \ell)$ columns of \tilde{E}_2^* to recover the symbols in the intersection of top m rows and bottom $(k - \ell)$ columns of \tilde{S}_2^* (depicted as (i) in Fig. 12). The symmetry of matrix S_2^* implies the decoding of the intersection of leftmost m columns and bottommost $(k - \ell)$ rows of \tilde{S}_2^* as well (depicted as (ii) in Fig. 12). We subtract these symbols out from the remaining data and update the values of $\{\tilde{E}_i^*, i = 1, 2, 3\}$ accordingly. Now, consider the rightmost $(k - 1 - \ell)$ columns of \tilde{E}_3^* . The rightmost $(k - 1 - \ell)$ columns of \tilde{S}_2^* are either zero or have been subtracted out, and as a result, the rightmost $(k - 1 - \ell)$ columns of \tilde{E}_3^* are identical to the $(k - 1 - \ell)$ columns of $\phi_\ell^* \tilde{S}_1^*$. Since the intersection of the bottom $(k - 1 - \ell)$ rows and the rightmost $(k - 1 - \ell)$ columns of \tilde{S}_1^* (depicted as (ix) in Fig. 12) comprises only of (zero-valued) message symbols, one can decode the rightmost $(k - 1 - \ell)$ columns of \tilde{S}_1^* from $\phi_\ell^* \tilde{S}_1^*$ (depicted as (iii) in Fig. 12). Since \tilde{S}_1^* is symmetric, this also implies decoding of the bottommost $(k - 1 - \ell)$ rows of \tilde{S}_1^* (depicted as (iv) in Fig. 12). Now assume that the values of these decoded symbols are subtracted out from $\{\tilde{E}_i^*, i = 1, 2, 3\}$. Once again we see that the ℓ th column of \tilde{S}_2^* comprises of symbols that are either zero-valued or have been subtracted out in the previous steps. As a result, the ℓ th column of \tilde{E}_3^* are identical to the ℓ th column of $\phi_\ell^* \tilde{S}_1^*$. Since the bottommost $(k - 1 - \ell)$ elements of the ℓ th column of \tilde{S}_1^* have been subtracted out, one can decode the ℓ th column of \tilde{S}_1^* from the ℓ th column of E_3^* (depicted as (v) in Fig. 12). Since \tilde{S}_1^* is symmetric, this also implies decoding of the ℓ th row of \tilde{S}_1^* (depicted as (vi) in Fig. 12). Now

assume that the values of these decoded symbols are also subtracted out from $\{\tilde{E}_i^*, i = 1, 2, 3\}$. We continue to focus on \tilde{E}_3^* and observe that the only non-zero elements remaining are the first $(\ell - 1)$ rows and columns. Obtaining the top-left $((\ell - 1) \times (\ell - 1))$ submatrices of \tilde{S}_1^* and \tilde{S}_2^* (depicted as (vii) and (viii) in Fig. 12) from this data is identical to the reconstruction requirement of an MSR product-matrix code with parameters $k = \ell$, $d = 2\ell - 2$ and $\alpha = \ell - 1$ in the absence of secrecy requirements. We already know [11] that this is possible, and this completes the proof of this step.

Step 2: We will now show that $H(\mathcal{E}) \leq R$. The eavesdropper obtains

$$\begin{aligned} E_1^* &= \Phi_m^* S_1^* \\ E_2^* &= \Phi_m^* S_2^* \\ E_3^* &= \Phi_\ell^* S_1^* + \Lambda_\ell \Phi_\ell^* S_2^*, \end{aligned}$$

where E_1^* and E_2^* are both sized $(m \times (k - 1))$, and E_3^* is sized $(\ell \times (k - 1))$. These three matrices together consist of

$$|\mathcal{E}| = 2m(k - 1) + \ell\alpha \quad (57)$$

symbols. Recall from (41) that $R = \ell\alpha + m(k - \ell)$. It thus suffices to show that least

$$|\mathcal{E}| - R = (2m(k - 1) + \ell\alpha) - (\ell\alpha + m(k - \ell)) \quad (58)$$

$$= m(k + \ell - 2) \quad (59)$$

of the symbols in $\{E_1^*, E_2^*, E_3^*\}$ are redundant.

To this end, first observe that the symmetry of matrices S_1^* and S_2^* makes the matrices $E_1^*(\Phi_m^*)^T$ and $E_2^*(\Phi_m^*)^T$ also symmetric. As a result, in each of the matrices E_1^* and E_2^* , $\frac{m(m-1)}{2}$ of the symbols are redundant. Now consider matrix E_3^* . The fact that Φ_m^* is a sub-matrix of Φ_ℓ^* implies that the matrix

$$E_1^* + \Lambda_m E_2^* = \Phi_m^* S_1^* + \Lambda_m \Phi_m^* S_2^*$$

comprises m rows of E_3^* . Thus the $m\alpha = m(k - 1)$ symbols in these m rows of E_3^* are redundant. Further, the relation

$$E_3^*(\Phi_m^*)^T = \Phi_\ell^*(E_1^*)^T + \Phi_\ell^*(E_2^*)^T \quad (60)$$

$$= \Phi_\ell^* S_1^*(\Phi_m^*)^T + \Phi_\ell^* S_2^*(\Phi_m^*)^T \quad (61)$$

implies that m symbols in the span of every row of E_3^* are redundant. Discounting the m rows of E_3^* that we previously showed to be redundant, we see that an additional $m(\ell - m)$ symbols of E_3^* are redundant. Adding these quantities, we get that a total of

$$2\frac{m(m-1)}{2} + m(k-1) + m(\ell-m) = m(k+\ell-2)$$

symbols in the data $\{E_1^*, E_2^*, E_3^*\}$ obtained by the eavesdropper are redundant, thus meeting the required condition (59).

Step 3: Having shown that $H(\mathcal{R}|\mathcal{E}, \mathcal{U}) = 0$ and $H(\mathcal{E}) \leq R$, the remainder of the proof is identical to the Step 3 of the proof of the MBR case (Theorem 4). \blacksquare

Proof of Theorem 11 ($\{\ell, m\}$ -security under the MSR code when $d > 2k - 2$): Consider an eavesdropper who gains access to the data stored in an arbitrary set of ℓ nodes, and in addition, to the data downloaded

for repair by some m of these ℓ nodes. We will now show that under code \mathcal{C}^* , this eavesdropper obtains no information pertaining to the message. To see this, imagine the existence of a set of i additional nodes that store only zeros as their data. Then, one can assume without loss of generality that the eavesdropper also has access to the data stored in these i nodes. Now, the resultant system is identical to one employing code \mathcal{C}'^* , with the eavesdropper having access to the data stored in a subset of $(\ell + i)$ nodes, and the data downloaded for repair by m of these $(\ell + i)$ nodes. As shown in Theorem 9, in this case, the eavesdropper obtains no information about the message in the system employing code \mathcal{C}'^* . Finally, since the two codes \mathcal{C}'^* and \mathcal{C}^* operate on the same message, it follows that the eavesdropper gains no information about the message in code \mathcal{C}^* . ■

REFERENCES

- [1] N. B. Shah, K. V. Rashmi, and P. V. Kumar, "Information-theoretically secure regenerating codes for distributed storage," in *Proc. Globecom*, Houston, Dec. 2011.
- [2] K. V. Rashmi, N. B. Shah, K. Ramchandran, and P. Kumar, "Regenerating codes for errors and erasures in distributed storage," in *Proc. International Symposium on Information Theory*, Jul. 2012.
- [3] D. Ford, F. Labelle, F. Popovici, M. Stokely, V. Truong, L. Barroso, C. Grimes, and S. Quinlan, "Availability in globally distributed storage systems," in *USENIX Symposium on Operating Systems Design and Implementation*, 2010.
- [4] K. V. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, "A solution to the network challenges of data recovery in erasure-coded distributed storage systems: A study on the Facebook warehouse cluster," in *Proc. USENIX HotStorage*, Jun. 2013.
- [5] S. Ghemawat, H. Gobioff, and S. Leung, "The Google file system," in *ACM SIGOPS Operating Systems Review*, vol. 37, no. 5. ACM, 2003, pp. 29–43.
- [6] D. Borthakur, "HDFS and Erasure Codes (HDFS-RAID). <http://hadoopblog.blogspot.com/2009/08/hdfs-and-erasure-codes-hdfs-raid.html>," Aug. 2009.
- [7] K. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, "A hitchhiker's guide to fast and efficient data reconstruction in erasure-coded data centers," in *Proceedings of the 2014 ACM conference on SIGCOMM*. ACM, 2014, pp. 331–342.
- [8] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "Xoring elephants: Novel erasure codes for big data," in *VLDB Endowment*, 2013.
- [9] K. V. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," in *Proc. 47th Annual Allerton Conference on Communication, Control, and Computing*, Urbana-Champaign, Sep. 2009, pp. 1243–1249.
- [10] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Explicit codes minimizing repair bandwidth for distributed storage," in *Proc. IEEE Information Theory Workshop (ITW)*, Cairo, Jan. 2010.
- [11] K. V. Rashmi, N. B. Shah, and P. V. Kumar, "Optimal exact-regenerating codes for the MSR and MBR points via a product-matrix construction," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5227–5239, Aug. 2011.
- [12] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Distributed storage codes with repair-by-transfer and non-achievability of interior points on the storage-bandwidth tradeoff," *IEEE Transactions on Information Theory*, vol. 58, no. 3, pp. 1837–1852, Mar. 2012.
- [13] —, "Interference alignment in regenerating codes for distributed storage: Necessity and code constructions," *IEEE Transactions on Information Theory*, vol. 58, no. 4, pp. 2134–2158, Apr. 2012.
- [14] D. Papailiopoulos, A. Dimakis, and V. Cadambe, "Repair optimal erasure codes through Hadamard designs," *IEEE Transactions on Information Theory*, vol. 59, no. 5, pp. 3021–3037, May 2013.
- [15] I. Tamo, Z. Wang, and J. Bruck, "Zigzag codes: MDS array codes with optimal rebuilding," *Information Theory, IEEE Transactions on*, vol. 59, no. 3, pp. 1597–1616, 2013.
- [16] V. Cadambe, C. Huang, J. Li, and S. Mehrotra, "Polynomial length MDS codes with optimal repair in distributed storage," in *Forty Fifth Asilomar Conference on Signals, Systems and Computers*, Nov. 2011, pp. 1850–1854.
- [17] Z. Wang, I. Tamo, and J. Bruck, "Long mds codes for optimal repair bandwidth," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2012, pp. 1182–1186.
- [18] C. Suh and K. Ramchandran, "Exact-repair MDS code construction using interference alignment," *IEEE Transactions on Information Theory*, pp. 1425–1442, Mar. 2011.
- [19] B. Sasidharan and P. V. Kumar, "High-rate regenerating codes through layering," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 1611–1615.

- [20] B. Gastón, J. Pujol, and M. Villanueva, "Quasi-cyclic flexible regenerating codes," *arXiv preprint arXiv:1209.3977*, 2012.
- [21] S.-J. Lin and W.-H. Chung, "Novel repair-by-transfer codes and systematic exact-mbr codes with lower complexities and smaller field sizes," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 25, no. 12, pp. 3232–3241, 2014.
- [22] J. Li, X. Tang, and P. Udaya, "A framework of constructions of minimum storage regenerating codes with the optimal update/access property for distributed storage systems based on invariant subspace technique," *arXiv preprint arXiv:1311.4947*, 2013.
- [23] M. Kurihara and H. Kuwakado, "Generalization of Rashmi-Shah-Kumar Minimum-Storage-Regenerating Codes," *arXiv preprint arXiv:1309.6701*, 2013.
- [24] H. Hou, K. W. Shum, M. Chen, and H. Li, "Basic regenerating code: Binary addition and shift for exact repair," in *IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2013, pp. 1621–1625.
- [25] M. Gerami, M. Xiao, M. Skoglund, K. W. Shum, and D. Lin, "Optimized-cost repair in multi-hop distributed storage systems with network coding," *arXiv preprint arXiv:1303.6046*, 2013.
- [26] C. Tian, V. Aggarwal, V. Vaishampayan *et al.*, "Exact-repair regenerating codes via layered erasure correction and block designs," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 1431–1435.
- [27] N. B. Shah, "On minimizing data-read and download for storage-node recovery," *IEEE Communications Letters*, 2013.
- [28] S. Jiekak and N. L. Scouarnec, "CROSS-MBCR: Exact minimum bandwidth coordinated regenerating codes," *arXiv preprint arXiv:1207.0854*, 2012.
- [29] "Citi data theft points up a nagging problem: CNBC," Jun. 2011. [Online]. Available: http://www.cnbc.com/id/43351159/Citi_Data_Theft_Points_Up_a_Nagging_Problem
- [30] "LinkedIn hack is latest blow to online confidence: USA Today," Jun. 2012. [Online]. Available: <http://usatoday30.usatoday.com/tech/news/story/2012-06-06/linkedin-investigates-security-breach/55423876/1>
- [31] "Yahoo confirms, apologizes for the email hack, says still fixing: TechCrunch," Jul. 2012. [Online]. Available: <http://techcrunch.com/2012/07/12/yahoo-confirms-apologizes-for-the-email-hack-says-still-fixing-plus-check-if-you-were-impacted-non-yahoo-accounts-apply/>
- [32] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4539–4551, 2010.
- [33] S. Pawar, S. El Rouayheb, and K. Ramchandran, "Securing dynamic distributed storage systems against eavesdropping and adversarial attacks," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6734–6753, 2011.
- [34] A. R. Lehman and E. Lehman, "Complexity classification of network information flow problems," in *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2004, pp. 142–150.
- [35] J. Feldman, T. Malkin, C. Stein, and R. Servedio, "On the capacity of secure network coding," in *Proc. 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.
- [36] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Médard, "Resilient network coding in the presence of Byzantine adversaries," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2007, pp. 616–624.
- [37] R. Koetter and F. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3579–3591, 2008.
- [38] D. Silva, F. R. Kschischang, and R. Koetter, "A rank-metric approach to error control in random network coding," *Information Theory, IEEE Transactions on*, vol. 54, no. 9, pp. 3951–3967, 2008.
- [39] C. Ngai and R. Yeung, "Secure error-correcting (SEC) network codes," in *IEEE Network Coding, Theory, and Applications (NetCod)*, 2009, pp. 98–103.
- [40] O. Kosut, L. Tong, and D. Tse, "Nonlinear network coding is necessary to combat general Byzantine attacks," in *Annual Allerton Conference on Communication, Control, and Computing*, 2009, pp. 593–599.
- [41] H. Yao, D. Silva, S. Jaggi, and M. Langberg, "Network codes resilient to jamming and eavesdropping," in *IEEE International Symposium on Network Coding (NetCod)*, 2010.
- [42] N. Cai and R. Yeung, "Secure network coding on a wiretap network," *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 424–435, Jan. 2011.
- [43] K. V. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit and optimal exact-regenerating codes for the minimum-bandwidth point in distributed storage," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Austin, Jun. 2010, pp. 1938–1942.
- [44] C. Tian, "Rate region of the (4, 3, 3) exact-repair regenerating codes," in *IEEE International Symposium on Information Theory*, Istanbul, Jul. 2013.
- [45] B. Sasidharan, K. Senthoo, and P. V. Kumar, "An improved outer bound on the storage-repair-bandwidth tradeoff of exact-repair regenerating codes," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 2430–2434.
- [46] S. Mohajer and R. Tandon, "New bounds on the (n, k, d) storage systems with exact repair," in *IEEE International Symposium on Information Theory (ISIT)*, 2015.
- [47] N. Prakash and M. N. Krishnan, "The storage-repair-bandwidth trade-off of exact repair linear regenerating codes for the case $d = k = n - 1$," in *IEEE International Symposium on Information Theory (ISIT)*, 2015.

- [48] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath, "Optimal locally repairable and secure codes for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 60, no. 1, pp. 212–236, 2014.
- [49] F. MacWilliams and N. Sloane, *The Theory of Error-Correcting Codes, Part I*. North-Holland Publishing Company, 1977.
- [50] A. Vulimiri, O. Michel, P. Godfrey, and S. Shenker, "More is less: Reducing latency via redundancy," in *11th ACM Workshop on Hot Topics in Networks*, Oct. 2012, pp. 13–18.
- [51] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica, "Why let resources idle? aggressive cloning of jobs with dolly," in *USENIX HotCloud*, Jun. 2012.
- [52] J. Dean and L. A. Barroso, "The tail at scale," *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [53] N. B. Shah, K. Lee, and K. Ramchandran, "When do redundant requests reduce latency?" in *Proc. 51st Annual Allerton Conference on Control, Computing, and Communication*, Urbana-Champaign, Oct. 2013.
- [54] G. Joshi, Y. Liu, and E. Soljanin, "On the delay-storage trade-off in content download from coded distributed storage systems," *Selected Areas in Communications, IEEE Journal on*, vol. 32, no. 5, pp. 989–997, 2014.
- [55] N. Le Scouarnec, "Exact scalar minimum storage coordinated regenerating codes," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2012, pp. 1197–1201.
- [56] Y. Hu, P. P. Lee, and K. W. Shum, "Analysis and construction of functional regenerating codes with uncoded repair for distributed storage systems," in *INFOCOM, 2013 Proceedings IEEE*. IEEE, 2013, pp. 2355–2363.
- [57] I. Tamo, Z. Wang, and J. Bruck, "MDS array codes with optimal rebuilding," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, St. Petersburg, Jul. 2011.
- [58] V. Cadambe, C. Huang, and J. Li, "Permutation code: optimal exact-repair of a single failed node in MDS code based distributed storage systems," in *IEEE International Symposium on Information Theory (ISIT)*, 2011, pp. 1225–1229.
- [59] B. Sasidharan, K. Senthoo, and P. V. Kumar, "An improved outer bound on the storage-repair-bandwidth tradeoff of exact-repair regenerating codes," *arXiv preprint arXiv:1312.6079*, 2013.
- [60] N. Silberstein, A. Rawat, and S. Vishwanath, "Adversarial error resilience in distributed storage using MRD codes and MDS array codes," *arXiv:1202.0800*, 2012.
- [61] S. Goparaju, S. El Rouayheb, R. Calderbank, and H. V. Poor, "Data secrecy in distributed storage systems under exact repair," in *Network Coding (NetCod), 2013 International Symposium on*. IEEE, 2013, pp. 1–6.
- [62] R. Tandon, S. Amuru, T. C. Clancy, and R. M. Buehrer, "Towards optimal secure distributed storage systems with exact repair," *arXiv preprint arXiv:1310.0054*, 2013.
- [63] T. Ernvall, S. El Rouayheb, C. Hollanti, and H. V. Poor, "Capacity and security of heterogeneous distributed storage systems," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 12, pp. 2701–2709, 2013.
- [64] Y. Han, R. Zheng, and W. Mow, "Exact regenerating codes for Byzantine fault tolerance in distributed storage," in *Proc. IEEE International Conference on Computer Communications (INFOCOM)*, Florida, USA, March 2012.
- [65] F. Oggier and A. Datta, "Byzantine fault tolerance of regenerating codes," in *IEEE International Conference on Peer-to-Peer Computing*, 2011, pp. 112–121.
- [66] K. W. Shum and Y. Hu, "Exact minimum-repair-bandwidth cooperative regenerating codes for distributed storage systems," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*. IEEE, 2011, pp. 1442–1446.
- [67] T. K. Dikaliotis, A. G. Dimakis, and T. Ho, "Security in distributed storage systems by communicating a logarithmic number of bits," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*. IEEE, 2010, pp. 1948–1952.
- [68] O. Kosut, "Polytope codes for distributed storage in the presence of an active omniscient adversary," in *IEEE International Symposium on Information Theory*, 2013, pp. 897–901.
- [69] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [70] L. Ozarow and A. Wyner, "Wire-tap channel II," in *Advances in cryptology*, 1985, pp. 33–50.
- [71] D. S. Bernstein, *Matrix mathematics: Theory, facts, and formulas with application to linear systems theory*. Princeton University Press, 2005.